Oracle® Grid Infrastructure Installation and Upgrade Guide





Oracle Grid Infrastructure Installation and Upgrade Guide, 21c for Microsoft Windows

F31798-04

Copyright © 2012, 2025, Oracle and/or its affiliates.

Primary Authors: Binika Kumar, Tejas N K

Contributing Authors: Douglas Williams, Prakash Jashnani, Subhash Chandra, Sunil Surabhi

Contributors: Alex Keh, Aneesh Khandelwal, Barb Glover, David Price, Deepak Keswani, Eric Belden, Esperanza Citlali Garcia Medina, Ian Cookson, James Williams, Jacqueline Sideri, Jiangqi Yang, Jonathan Creighton, Joseph Francis, Kevin Jernigan, Luv Sharma, Malai Stalin, Maria Colgan, Michael Coulter, Pallavi Kamath, Parvathi Subramanian, Prasad Bagal, Priyanka Sharma, Rajesh Dasari, Ramesh Chakravarthula, Richard Wessman, Robert Achacoso, Rudregowda Mallegowda, Sivaselvam Narayanasamy, Sreyasi Ghoshal, Srinivas Poovala, Subrahmanyam Kodavaluru, Susheel Chauhan, Suresh Yambari Venkata Naga, Tonghua Li, Vishal Saxena, Yanhua Xia

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

		•	
D	$r \alpha$	ta.	ce

	Audience	
	Documentation Accessibility	j
	Diversity and Inclusion	ii
	Set Up Java Access Bridge to Implement Java Accessibility	ii
	Conventions	i
1	Oracle Grid Infrastructure Installation Checklist	
	Oracle Grid Infrastructure Installation Server Hardware Checklist	1
	Operating System Checklist for Oracle Grid Infrastructure and Oracle RAC	2
	Server Configuration Checklist for Oracle Grid Infrastructure	3
	Oracle Grid Infrastructure Network Checklist	4
	User Environment Configuration Checklist for Oracle Grid Infrastructure	5
	Storage Checklist for Oracle Grid Infrastructure	6
	Installer Planning Checklist for Oracle Grid Infrastructure	7
2	Configuring Servers for Oracle Grid Infrastructure and Oracle RA Checking Server Hardware and Memory Configuration	C 1
	Checking the Available RAM on Windows Systems	2
	Checking the Currently Configured Virtual Memory on Windows Systems	2
	Checking the System Processor Type	2
	Checking the Available Disk Space for Oracle Home Directories	2
	Checking the Available TEMP Disk Space	3
	General Server Requirements	4
	Server Minimum Hardware and Memory Requirements	4
	Server Minimum Storage Requirements	5
	Disk Format Requirements	6
	Configuring Time Synchronization for the Cluster	6
	About Cluster Time Synchronization	6
	Understanding Network Time Requirements	7
	Configuring the Windows Time Service	7
	Configuring Network Time Protocol	8

3 Configuring Operating Systems for Oracle Grid Infrastructure and Oracle RAC

Reviewing Operating System and Software Upgrade Best Practices	1
General Upgrade Best Practices	2
Oracle ASM Upgrade Notifications	2
Rolling Upgrade Procedure Notifications	2
Reviewing Operating System Security Common Practices	3
Verifying Privileges for Copying Files in the Cluster	3
Identifying Software Requirements	3
Windows Firewall Feature on Windows Servers	4
Checking the Operating System Version	5
Checking Hardware and Software Certification on My Oracle Support	5
Installation Requirements for Web Browsers	6
Configuring Networks for Oracle Grid Infrastructure and Ora	acle RAC
About Oracle Grid Infrastructure Network Configuration Options	2
Understanding Network Addresses	2
About the Public IP Address	3
About the Private IP Address	3
About the Virtual IP Address	3
About the Grid Naming Service (GNS) Virtual IP Address	4
About the SCAN	4
Network Interface Hardware Requirements	5
Network Requirements for Each Node	6
Network Requirements for the Private Network	6
Network Requirements for the Public Network	7
IPv6 Protocol Support for Windows	7
Using Multiple Public Network Adapters	g
Network Interface Configuration Options for Performance	g
Network Configuration Tasks for Windows Server Deployments	g
Disabling Windows Media Sensing	10
Deconfigure DNS Registration for Public Network Adapter	10
Manually Configure Automatic Metric Values	11
Setting UDP and TCP Dynamic Port Range for Oracle RAC Installations	11
Oracle Grid Infrastructure IP Name and Address Requirements	11
About Oracle Grid Infrastructure Name Resolution Options	12
Cluster Name and SCAN Requirements	13
IP Name and Address Requirements For Grid Naming Service (GNS)	14

4

	IP Name and Address Requirements For Multi-Cluster GNS	14
	About Multi-Cluster GNS Networks	14
	Configuring GNS Server Clusters	15
	Configuring GNS Client Clusters	15
	Creating and Using a GNS Client Data File	15
	IP Address Requirements for Manual Configuration	16
	Confirming the DNS Configuration for SCAN	17
	Grid Naming Service Cluster Configuration Example	17
	Domain Delegation to Grid Naming Service	18
	Choosing a Subdomain Name for Use with Grid Naming Service	19
	Configuring DNS for Domain Delegation to Grid Naming Service	19
	Manual IP Address Configuration Example	20
	Intended Use of Network Adapters	21
	Broadcast Requirements for Networks Used by Oracle Grid Infrastructure	22
	Multicast Requirements for Networks Used by Oracle Grid Infrastructure	22
	Configuring Multiple ASM Interconnects on Microsoft Windows Platforms	22
5	Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle RAC	
	Creating Installation Groups and Users for Oracle Grid Infrastructure and Oracle RAC	1
	About the Oracle Installation User	2
	About the Oracle Home User for the Oracle RAC Installation	2
	When to Create an Oracle Home User	3
	Restrictions and Guidelines for Oracle Home Users	4
	Determining if an Oracle Home User Exists	4
	Creating an Oracle Home User	4
	Using an Existing Oracle Software Owner User	5
	Oracle Home User Configurations for Oracle Installations	5
	Understanding the Oracle Inventory Directory and the Oracle Inventory Group	5
	Standard Administration and Job Role Separation User Groups About Job Role Separation Operating System Privileges Groups and Hears	6
	About Job Role Separation Operating System Privileges Groups and Users Oracle Software Owner for Each Oracle Software Product	6 7
	Standard Oracle Database Groups for Database Administrators	8
	Oracle ASM Groups for Job Role Separation	9
	Extended Oracle Database Administration Groups for Job Role Separation	9
	Operating System Groups Created During Installation	10
	Example of Using Role-Allocated Groups and Users	13
	Configuring User Accounts	14
	Configuring Environment Variables for the Oracle Installation User	15
	Verifying User Privileges to Update Remote Nodes	15
	Managing User Accounts with User Account Control	16
	managing Oser Accounts with Oser Account Control	10

	Creating Oracle Software Directories	16
	Temporary Directories	17
	Oracle Base Directory for the Grid User	17
	Oracle Home Directory for Grid Infrastructure Management Repository	19
	Grid Home Directory	19
	Oracle Inventory Directory	20
	Requirements for the Oracle Grid Infrastructure Home Directory	20
	About Creating the Oracle Base Directory Path	21
6	Configuring Shared Storage for Oracle Database and Grid Infra	structure
	Supported Storage Options for Oracle Grid Infrastructure and Oracle RAC	2
	Storage Considerations for Oracle Grid Infrastructure and Oracle RAC	3
	Guidelines for Choosing a Shared Storage Option	3
	Guidelines for Using Oracle ASM Disk Groups for Storage	4
	Guidelines for Using Direct Network File System (NFS) with Oracle RAC	4
	About Migrating Existing Oracle ASM Instances	5
	Preliminary Shared Disk Preparation	5
	Disabling Write Caching	5
	Enabling Automounting for Windows	6
	Configuring Disk Partitions on Shared Storage	7
	Creating Disk Partitions Using the Disk Management Interface	7
	Creating Disk Partitions using the DiskPart Utility	7
7	Configuring Storage for Oracle Automatic Storage Managemen	t
	Identifying Storage Requirements for Using Oracle ASM for Shared Storage	1
	Oracle Clusterware Storage Space Requirements	6
	About the Grid Infrastructure Management Repository	7
	Guidelines for the Grid Infrastructure Management Repository	7
	Restrictions for Disk Partitions Used By Oracle ASM	8
	Preparing Your System to Use Oracle ASM for Shared Storage	8
	Identifying and Using Existing Oracle Database Disk Groups on Oracle ASM	8
	Selecting Disks to use with Oracle ASM Disk Groups	g
	Specifying the Oracle ASM Disk Discovery String	g
	Marking Disk Partitions for Oracle ASM Before Installation	10
	Using asmtoolg to Mark Disks	10
	Using asmtoolg to Remove Disk Stamps	11
	asmtool Command Line Reference	11

8 Configuring Direct NFS Client for Oracle RAC Data Files

	About Direct NFS Client Storage	1
	Creating an oranfstab File for Direct NFS Client	3
	Configurable Attributes for the oranfstab File	5
	Mounting NFS Storage Devices with Direct NFS Client	7
	Specifying Network Paths for a NFS Server	7
	Enabling Direct NFS Client	8
	Performing Basic File Operations Using the ORADNFS Utility	8
	Monitoring Direct NFS Client Usage	8
	Disabling Direct NFS Client	9
9	Installing Oracle Grid Infrastructure for a Cluster	
	About Image-Based Oracle Grid Infrastructure Installation	2
	Setup Wizard Installation Options for Creating Images	3
	Understanding Cluster Configuration Options	3
	About Oracle Standalone Clusters	4
	About Oracle Extended Clusters	4
	About Default File Permissions Set by Oracle Universal Installer	5
	Installing Oracle Grid Infrastructure for a New Cluster	5
	Installing Oracle Grid Infrastructure Using a Cluster Configuration File	8
	Installing Software Binaries for Oracle Grid Infrastructure for a Cluster	10
	Configuring the Software Binaries for Oracle Grid Infrastructure for a Cluster	11
	Configuring the Software Binaries Using a Response File	12
	Setting Ping Targets for Network Checks	12
	Confirming Oracle Clusterware Function	13
	Confirming Oracle ASM Function for Oracle Clusterware Files	14
	Understanding Offline Processes in Oracle Grid Infrastructure	14
10	Oracle Grid Infrastructure Postinstallation Tasks	
	Required Postinstallation Tasks	1
	Downloading Release Update Patches	2
	Configuring Exceptions for the Windows Firewall	2
	About Configuring Exceptions for the Windows Firewall	3
	Firewall Exceptions for Oracle Database	4
	Firewall Exceptions for Oracle Database Examples (or the Companion CD)	4
	Firewall Exceptions for Oracle Gateways	4
	Firewall Exceptions for Oracle Clusterware and Oracle ASM	5
	Firewall Exceptions for Oracle RAC Database	6
	Firewall Exceptions for Other Oracle Products	6

	Troubleshooting Windows Firewall Exceptions	6
	Recommended Postinstallation Tasks	7
	About Installing Oracle Autonomous Health Framework	7
	Optimize Memory Usage for Programs	8
	Create a Fast Recovery Area Disk Group	8
	About the Fast Recovery Area and the Fast Recovery Area Disk Group	8
	Creating the Fast Recovery Area Disk Group	9
	Checking the SCAN Configuration	10
	Installing Grid Infrastructure Management Repository	10
	About Grid Infrastructure Management Repository Installation	11
	Creating the Grid Infrastructure Management Repository Disk Group	11
	Installing Grid Infrastructure Management Repository for a New Cluster	12
	Adding Grid Infrastructure Management Repository to an Existing Cluster	13
	Using Earlier Oracle Database Releases with Grid Infrastructure	15
	General Restrictions for Using Earlier Oracle Database Releases	15
	Using the Correct LSNRCTL Commands	15
	Modifying Oracle Clusterware Binaries After Installation	15
ΤТ	Upgrading Oracle Grid Infrastructure	
	Understanding Out-of-Place and Rolling Upgrades	2
	About Oracle Grid Infrastructure Upgrade and Downgrade	3
	Options for Oracle Grid Infrastructure Upgrades	3
	Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades	4
	Preparing to Upgrade an Existing Oracle Clusterware Installation	5
	Upgrade Checklist for Oracle Grid Infrastructure	6
	Tasks to Complete Before Upgrading Oracle Grid Infrastructure	7
	Create an Oracle ASM Password File	9
	Running the Oracle ORAchk Upgrade Readiness Assessment	9
	Using CVU to Validate Readiness for Oracle Clusterware Upgrades	10
	About the CVU Grid Upgrade Validation Command Options	10
	Example of Verifying System Upgrade Readiness for Grid Infrastructure	11
	Understanding Rolling Upgrades Using Batches	11
	Performing Rolling Upgrades of Oracle Grid Infrastructure	12
	Upgrading Oracle Grid Infrastructure from an Earlier Release	12
	Completing an Oracle Clusterware Upgrade when Nodes Become Unreachable	14
	Joining Inaccessible Nodes After Forcing an Upgrade	15
	Changing the First Node for Install and Upgrade	15
	Upgrading Grid Infrastructure Management Repository from an Earlier Release	16
	Applying Patches to Oracle Grid Infrastructure	17
	About Individual Oracle Grid Infrastructure Patches	17
	About Oracle Grid Infrastructure Software Patch Levels	17

Patching Oracle Grid Infrastructure	18
Applying Patches to Grid Infrastructure Management Repository	18
In-Place Grid Infrastructure Management Repository Patching	19
Out-of-Place Grid Infrastructure Management Repository Patching	19
In-Place Grid Infrastructure Management Repository Patch Rollback	20
Out-of-Place Grid Infrastructure Management Repository Patch Rollback	21
Updating Oracle Enterprise Manager Cloud Control Target Parameters	21
Updating the Enterprise Manager Cloud Control Target After Upgrades	22
Updating the Enterprise Manager Agent Base Directory After Upgrades	22
Registering Resources with Oracle Enterprise Manager After Upgrades	22
Checking Cluster Health Monitor Repository Size After Upgrading	23
Downgrading Oracle Clusterware to an Earlier Release	24
About Oracle Grid Infrastructure Downgrades	24
Options for Oracle Grid Infrastructure Downgrades	25
Restrictions for Oracle Grid Infrastructure Downgrades	25
Downgrading Oracle Standalone Cluster	25
Downgrading Oracle Grid Infrastructure when Upgrade Fails	27
Completing Failed or Interrupted Installations and Upgrades	27
Completing Failed Installations and Upgrades	28
Continuing Incomplete Upgrade of First Nodes	28
Continuing Incomplete Upgrades on Remote Nodes	29
Continuing Incomplete Installations on First Nodes	29
Continuing Incomplete Installation on Remote Nodes	30
Modifying or Deinstalling Oracle Grid Infrastructure	
Deciding When to Deinstall Oracle Clusterware	1
Migrating Standalone Grid Infrastructure Servers to a Cluster	2
Changing the Oracle Grid Infrastructure Home Path	3
Unconfiguring Oracle Clusterware Without Removing the Software	5
Removing Oracle Clusterware and Oracle ASM Software	6
About Oracle Deinstallation Options	6
Files Deleted by the deinstall Command	7
deinstall Command Reference	8
Using the deinstall Command to Remove Oracle Clusterware and Oracle ASM	10
Running the deinstall Command From an Oracle Home	10
Running the deinstall command with a Response File	11
Generating a Response File For Use With the deinstall Command	11

A Installing and Configuring Oracle Grid Infrastructure Using Response Files

How Response Files Work	A-1
Deciding to Use Silent Mode or Response File Mode	A-2
Using Response Files	A-3
Preparing Response Files	A-3
About Response File Templates	A-3
Editing a Response File Template	A-4
Recording Response Files	A-5
Running Oracle Universal Installer Using a Response File	A-6
Running Oracle Net Configuration Assistant Using Response Files	A-7
Postinstallation Configuration Using Response File Created During Installation	A-8
Using the Installation Response File for Postinstallation Configuration	A-8
Running Postinstallation Configuration Using a Response File	A-9
Postinstallation Configuration Using the ConfigToolAllCommands Script	A-10
About the Postinstallation Configuration File	A-11
Creating a Password Response File	A-11
Running Postinstallation Configuration Using a Script and Response Files	A-12
Optimal Flexible Architecture	
About the Optimal Flexible Architecture Standard	B-1
About Multiple Oracle Homes Support	B-2
Oracle Base Directory Naming Convention	B-2
Oracle Home Directory Naming Convention	B-3
Ontimal Elexible Architecture File Path Examples	R-4

В



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).

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

- Audience
- Documentation Accessibility
- · Diversity and Inclusion
- Set Up Java Access Bridge to Implement Java Accessibility
 Install Java Access Bridge so that assistive technologies on Microsoft Windows systems can use the Java Accessibility API.
- Conventions

Audience

Oracle Grid Infrastructure Installation and Upgrade Guide for Microsoft Windows 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 customers with specialized system roles who intend to install Oracle Real Application Clusters (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 Administrator user privileges. When the Oracle Grid Infrastructure installation and configuration is successfully completed, a system administrator should only provide configuration information and grant access to the database administrator to run scripts that require Administrator user privileges during an Oracle RAC installation.

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

Documentation Accessibility

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

Access to Oracle Support

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



Diversity and Inclusion

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

Set Up Java Access Bridge to Implement Java Accessibility

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

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

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

Related Topics

Java Platform, Standard Edition Java Accessibility Guide

Conventions

The following text conventions are used in this document:

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

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.

- Oracle Grid Infrastructure Installation Server Hardware Checklist
 Review server hardware requirements for Oracle Grid Infrastructure installation.
- Operating System Checklist for Oracle Grid Infrastructure and Oracle RAC Review the following operating system checklist for all installations.
- <u>Server Configuration Checklist for Oracle Grid Infrastructure</u>
 Use this checklist to check minimum server configuration requirements for Oracle Grid Infrastructure installations.
- Oracle Grid Infrastructure Network Checklist
 Review the installation requirements to ensure that you have the required hardware, names, and addresses for the cluster.
- <u>User Environment Configuration Checklist for Oracle Grid Infrastructure</u>
 Review the following environment checklist for all installations.
- Storage Checklist for Oracle Grid Infrastructure
 Review the checklist for storage hardware and configuration requirements for Oracle Grid Infrastructure installation.
- Installer Planning Checklist for Oracle Grid Infrastructure
 Review the checklist for planning your Oracle Grid Infrastructure installation before starting
 Oracle Universal Installer.

Oracle Grid Infrastructure Installation Server Hardware Checklist

Review server hardware requirements for Oracle Grid Infrastructure installation.

Table 1-1 Server Hardware Checklist for Oracle Grid Infrastructure

Check	Task
Server Display Cards	At least 1024 x 768 display resolution for Oracle Universal Installer. Confirm display monitor.
Minimum Random Access Memory (RAM)	At least 16 GB RAM for Oracle Grid Infrastructure installations, including installations where you plan to install Oracle RAC.
Virtual Memory configuration	Oracle recommends that you set the paging file size to match the amount of RAM available on the server, up to a maximum of 16 GB



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

Check	Task
Intelligent Platform Management Interface (IPMI) Configuration	 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 user name and password to provide when prompted during installation. i) Note IPMI is not available on Microsoft Windows for Oracle Database Release 21c, but will be available in a future release.

Related Topics

Configuring Servers for Oracle Grid Infrastructure and Oracle RAC
 You must complete certain operating system tasks on your servers before you install
 Oracle Grid Infrastructure for a Cluster and Oracle Real Application Clusters.

Operating System Checklist for Oracle Grid Infrastructure and Oracle RAC

Review the following operating system checklist for all installations.

Table 1-2 Operating System General Checklist for Oracle Grid Infrastructure and Oracle RAC

Check	Task
Review operating system security standards	Secure operating systems are an important basis for general system security.
Verify Privileges for Copying Files in the Cluster	During installation, OUI copies the software from the local node to the remote nodes in the cluster. The installation user must have Administrator privileges on the other nodes in the cluster to copy the files.
Configure Remote Desktop Services	If you want to enable remote display of the installation process, then configure remote access for necessary user accounts.
Disable Windows Firewall	If the Windows Firewall is enabled, then remote copy and configuration assistants such as Virtual IP Configuration Assistant (VIPCA), Network Configuration Assistant (NETCA), and Oracle Database Configuration Assistant (DBCA) will fail during Oracle RAC installation.



Table 1-2 (Cont.) Operating System General Checklist for Oracle Grid Infrastructure and Oracle RAC

Check	Task	
Operating system	Review the system requirements section for a list of minimum software component versions. Use the same operating system version on each cluster member node.	
general requirements	 Windows Server 2022 x64 - Standard, Datacenter, and Essentials editions (Supported starting Oracle Database 21c Release Update 21.4 or later) Windows Server 2019 x64 - Standard, Datacenter, and Essentials editions Windows Server 2016 x64 - Standard, Datacenter, and Essentials editions Windows Server 2012 R2 x64 - Standard, Datacenter, Essentials, and Foundation editions 	
Ensure web browsers are supported by Enterprise Manager	Web browsers must support Java Script, and the HTML 4.0 and CSS 1.0 standards.	

Server Configuration Checklist for Oracle Grid Infrastructure

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

Table 1-3 Server Configuration Checklist for Oracle Grid Infrastructure

Check	Task
Disk space allocated to the temporary file system	At least 1 GB of space in the temporary directory. Oracle recommends 2 GB or more.
Storage hardware for Oracle software	Attached storage hardware can be either Storage Area Network (SAN) storage or Network-Attached Storage (NAS).
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. Additionally, the file path names cannot contain spaces.
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), Oracle Cluster Time Synchronization Service, or Windows Time Service.



Oracle Grid Infrastructure Network Checklist

Review the installation requirements to ensure that you have the 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, and not available for Oracle Grid Infrastructure use.

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; use NIC teaming so they appear as one Ethernet port name.
	 The switches and network interface cards must be at least 1 GbE.
	The network protocols are UDP and TCP/IP.
Private Network Hardware for the Interconnect	 Private dedicated network switch (redundant switches recommended, with NIC teaming so they appear as one Ethernet port name), connected to the private interface ports for each cluster member node.
	 The switches and network interface adapters must be at least 1 GbE, with 10 GbE recommended. Alternatively, use InfiniBand for the interconnect.
	 The interconnect must support the user datagram protocol (UDP).
Perform Windows- specific network configuration tasks	Disable the Media Sensing Feature for TCP/IP. Media Sense allows Windows to uncouple an IP address from a network interface card when the link to the local switch is lost.
·	• Deselect Automatic Registration with DNS for the Public Network Interface. To prevent Windows Server from potentially registering the wrong IP addresses for the node in DNS after a server restart, you must deconfigure the "Register this connection's addresses in DNS" option for the public network adapters.
	 Manually configure automatic metric values. Automatic Metric feature automatically configures the metric for the local routes that are based on link speed. If you use the default values for this feature, OUI sometimes selects the private network interface as the default public host name for the server when installing Oracle Grid Infrastructure.
	You must manually configure metric values for the public and private network interfaces to ensure that your public network adapter is first in the binding order and the private network adapter is second. To do so, set the public network interface metric to a lower value than the private network interface. This assigns a higher priority to the public network in the binding order.
Oracle Flex ASM Network Hardware	Oracle Flex ASM can use either the same private networks as Oracle Clusterware, or use its own dedicated private networks. Each network can be classified PUBLIC or PRIVATE+ASM or PRIVATE or ASM. Oracle ASM networks use the TCP protocol.



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

Check Task

Cluster Names and Addresses

Determine and configure the following names and addresses for the cluster:

- Cluster name: Decide a name for the cluster, and be prepared to enter it during installation. The cluster name should have the following characteristics:
 - Globally unique across all hosts, even across different DNS domains.
 - Must be between 1 and 15 characters in length.
 - 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).
- Grid Naming Service Virtual IP Address (GNS VIP): If you plan to use GNS, then configure a GNS name and fixed address on the DNS for 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. SCANs 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).

Cluster Node Public, Private and Virtual IP names and Addresses If you are configuring a cluster with Grid Naming Service (GNS), then OUI displays the public and virtual host name addresses labeled as "AUTO" because they are configured automatically. To use this option, you must have configured a subdomain on your DNS that is delegated to GNS for resolution, and you must have a fixed GNS VIP address where the delegated service requests can be routed.

If you are not using GNS, then configure the following for each cluster node:

- Public node name and address, configured on the DNS and in the hosts (for example, node1.example.com, address 192.0.2.10). The public node name should be the primary host name of each node, which is the name displayed by the hostname command.
- Private node address, configured on the private interface for each node. The installer identifies addresses in the private range as private IP addresses by default. For example: 10.0.0.10

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, node1vip.example.com, address 192.0.2.11)

If you are not using dynamic networks with GNS and subdomain delegation, then determine a virtual host name for each node. A virtual host name is a public node name that is used to reroute client requests sent to the node if the node is down. Oracle Database uses VIPs for client-to-database connections, so the VIP address must be publicly accessible. Oracle recommends that you provide a name in the format *hostname*-vip. For example: myclstr2-vip.

User Environment Configuration Checklist for Oracle Grid Infrastructure

Review the following environment checklist for all installations.



Table 1-5 Environment Configuration for Oracle Grid Infrastructure and Oracle RAC

Check	Task
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 install guide. Oracle Installation users have different requirements from Oracle Home users. User names must use only ASCII characters.
Configure the Oracle Software Owner Environment	Set the TEMP environment variable.
Unset Oracle software environment variables	If you have set ORA_CRS_HOME as an environment variable, then unset it before starting an installation or upgrade. Do not use ORA_CRS_HOME as a user environment variable.
	If you have had an existing installation on your system, and you are using the same user account to install this installation, then unset the ORACLE_HOME, ORACLE_BASE, ORACLE_SID, and TNS_ADMIN environment variables. Unset any other environment variable or system variable set for the Oracle installation user that is connected with Oracle software homes, such as ORA_NLS10.
	For example, if you have set ORACLE_HOME as a system environment variable with Oracle RAC, then it could lead to installation errors even if the value is null.
Manage User Account Control	If you have enabled the User Account Control security feature, then Oracle Universal Installer prompts you for either your consent or your credentials when installing Oracle Database software. Provide either the consent or your Windows Administrator credentials as appropriate.

Related Topics

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

You must configure certain users, groups, and environment settings used during Oracle Grid Infrastructure for a Cluster and Oracle Real Application Clusters installations.

Storage Checklist for Oracle Grid Infrastructure

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

Table 1-6 Oracle Grid Infrastructure Storage Configuration Checks

Check	Task
Minimum disk space (local or shared) for Oracle Grid Infrastructure Software	 At least 12 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home). Oracle recommends that you allocate 100 GB to allow additional space for patches. At least 3.5 GB of space for the Oracle base of the Oracle Grid Infrastructure installation owner (Grid user). The Oracle base includes Oracle Clusterware and Oracle ASM log files. At least 10 GB of space for Oracle Database Enterprise Edition. Allocate additional storage space as per your cluster configuration, as described in Oracle Clusterware Storage Space Requirements.



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 Oracle ASM on NFS, is locally mounted on each of the cluster nodes.
	Voting files are files that Oracle Clusterware uses to verify cluster node membership and status. Oracle Cluster Registry files (OCR) contain cluster and database configuration information for Oracle Clusterware.
Select Grid Infrastructure Management Repository	For Oracle Standalone Cluster deployment, you can specify the same or separate Oracle ASM disk group for the GIMR.
(GIMR) Storage Option	① Note
	Starting with Oracle Grid Infrastructure 19c, configuring GIMR is optional for Oracle Standalone Cluster deployments.

Related Topics

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, and installing the starter database, for various Oracle Cluster deployments.

Installer Planning Checklist for Oracle Grid Infrastructure

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

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

Check	Task
Read the Release Notes	See Oracle Database Release Notes.
Review the Licensing Information	You are permitted to use only those components in the Oracle Database for which you have purchased licenses. See <i>Oracle Database Licensing Information User Manual</i> .



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

Check	Task
Review Oracle Support Certification Matrix	New platforms and operating system software versions can 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/
	You must register online before using My Oracle Support. After logging in, from the menu options, select the Certifications tab. On the Certifications page, use the Certification Search options to search by Product, Release, and Platform. You can also search using the Certification Quick Link options such as Product Delivery, and Lifetime Support.
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. See <i>Oracle Clusterware Administration and Deployment Guide</i> .
Download and run Oracle ORAchk for runtime and upgrade checks, or runtime health checks	The Oracle ORAchk utility provides system checks that can help to prevent issues after installation. These checks include kernel requirements, operating system resource allocations, and other system requirements.
	Use the Oracle ORAchk Upgrade Readiness Assessment to obtain an automated upgrade-specific system health check for upgrades. For example:
	./orachk -u -o pre
	The Oracle ORAchk Upgrade Readiness Assessment automates many of the manual pre- and post-upgrade checks described in Oracle upgrade documentation.
	https://support.oracle.com/rs?type=doc&id=1268927.1
Ensure Task Scheduler jobs do not run during installation	If the installer is running when daily scheduled jobs start, then you may encounter unexplained installation problems if your scheduled job is performing cleanup, and temporary files are deleted before the installation is finished. Oracle recommends that you complete installation before daily scheduled jobs are run, or disable daily scheduled 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 an Oracle RAC Database node: This installation requires an upgrade of Oracle Grid Infrastructure, as Oracle Grid Infrastructure 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.



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

Check	Task
Unzip utility	Unzip 6.0 or later.
	Unzip is required to extract the image files for Oracle Database and Oracle Grid Infrastructure installations.

Configuring Servers for Oracle Grid Infrastructure and Oracle RAC

You must complete certain operating system tasks on your servers before you install Oracle Grid Infrastructure for a Cluster and Oracle Real Application Clusters.

The values provided in this chapter are for installation minimums only. Oracle recommends that you configure production systems in accordance with planned system loads.

- Checking Server Hardware and Memory Configuration
 Perform these tasks to gather your current system information.
- General Server Requirements
 Verify that servers where you install Oracle Grid Infrastructure meet the minimum
- requirements for installation.
 Server Minimum Hardware and Memory Requirements
 Each system must meet certain minimum hardware and memory requirements.
- <u>Server Minimum Storage Requirements</u> Each system must meet certain minimum storage requirements.
- <u>Disk Format Requirements</u>
 Oracle recommends that you install Oracle software, or binaries, on New Technology File System (NTFS) formatted drives or partitions.
- Configuring Time Synchronization for the Cluster
 Oracle Clusterware requires the same time zone setting on all cluster nodes.

Related Topics

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

Checking Server Hardware and Memory Configuration

Perform these tasks to gather your current system information.

- Checking the Available RAM on Windows Systems
 Use the control panel to check the available RAM on each server.
- Checking the Currently Configured Virtual Memory on Windows Systems
 Virtual memory (also known as a paging file) stores information that cannot fit in RAM, the main memory for the computer. All processes share the paging files, and a lack of space in the paging files can prevent processes from allocating memory.
- <u>Checking the System Processor Type</u>
 To view your processor type (32-bit or 64-bit), perform these steps.
- Checking the Available Disk Space for Oracle Home Directories
 Additional shared disk space on a cluster file system is required for the Oracle Grid
 Infrastructure Management Repository (GIMR), Oracle Cluster Registry (OCR) and voting
 files used by Oracle Clusterware.



Checking the Available TEMP Disk Space

The amount of disk space available in the TEMP directory is equivalent to the total amount of free disk space, minus what is needed to install the Oracle software.

Checking the Available RAM on Windows Systems

Use the control panel to check the available RAM on each server.

The minimum required RAM is 16 GB for Oracle Grid Infrastructure for Cluster installations, including installations where you plan to install Oracle RAC.

To determine the physical RAM size, for a computer, you can use either of the following methods:

- Open System in the control panel and Select the General tab.
- Alternatively, start the Windows Task Manager, then select the **Performance** tab to view the available memory for your system.

Checking the Currently Configured Virtual Memory on Windows Systems

Virtual memory (also known as a paging file) stores information that cannot fit in RAM, the main memory for the computer. All processes share the paging files, and a lack of space in the paging files can prevent processes from allocating memory.

Oracle recommends that you set the paging file size to match the amount of RAM available on the server, up to a maximum of 16 GB. If possible, split the paging file into multiple files on multiple physical devices. This configuration encourages parallel access to virtual memory, and improves the software performance.

- 1. From the Control panel, select System.
- 2. In the System Properties window, select the **Advanced** tab.
- 3. Under Performance, click **Performance Options**, or **Settings**.
- 4. In the Performance Options window, click the Advanced tab.

The virtual memory configuration is displayed at the bottom of the window. If necessary, refer to your operating system documentation for information about how to configure additional virtual memory.

Checking the System Processor Type

To view your processor type (32-bit or 64-bit), perform these steps.

- 1. From the Start menu, select **Run**. In the Run window, type in msinfo32.exe.
- 2. In the System Summary display, locate the System Type entry.
 - If the value for System Type is x64-based PC, then you have a 64-bit system.
 - If the value for System Type is x86-based PC, then you have a 32-bit system.

Checking the Available Disk Space for Oracle Home Directories

Additional shared disk space on a cluster file system is required for the Oracle Grid Infrastructure Management Repository (GIMR), Oracle Cluster Registry (OCR) and voting files used by Oracle Clusterware.

If you are installing Oracle RAC, then you must configure additional disk space for:



- The Oracle RAC software and log files
- The shared data files and, optionally, the shared Fast Recovery Area in an Oracle ASM disk group

If you use standard redundancy for Oracle Clusterware files, which is 3 Oracle Cluster Registry (OCR) files and 3 voting files, then you should have at least 2 GB of disk space available on three separate physical disks reserved for storing the Oracle Clusterware files in Oracle ASM.

(i) Note

You cannot create OCR or voting files (Oracle Clusterware files) on raw partitions. You can store Oracle Clusterware files only on Oracle ASM. Raw devices can be configured as Oracle ASM disks.

To ensure high availability of OCR or voting files on Oracle ASM, you need to have at least 2 GB of disk space for Oracle Clusterware files in three separate failure groups, with at least three physical disks. Each disk must have at least 1 GB of capacity to ensure that there is sufficient space to create Oracle Clusterware files.

If the temp space and the Grid home are on the same file system, then add together their respective requirements for the total minimum space required for that file system.

(i) Note

Oracle recommends that you choose the Oracle Grid Infrastructure Management Repository option when installing Oracle Grid Infrastructure. When you choose this option, OUI configures a Oracle Grid Infrastructure Management Repository database on one of the nodes in the cluster.

To determine the amount of available free disk space, there are two methods you can use:

- Using the Computer properties window:
 - a. Open the Start menu, then click Computer.
 - b. View the free disk space for each drive.
 - Right-click the drive on which you plan to install the Oracle software and select **Properties** to view additional information about the disk space for that drive.
- Using the Disk Management Microsoft Management Console (MMC) plug-in:
 - From the **Start** menu, select **Run..**.
 - In the Run window, type in Diskmgmt.msc to open the Disk Management graphical user interface (GUI).

The Disk Management GUI displays the available space on the available file systems.

Checking the Available TEMP Disk Space

The amount of disk space available in the TEMP directory is equivalent to the total amount of free disk space, minus what is needed to install the Oracle software.

Log in as or switch to the user that will be performing the installation. See "About the Oracle Installation User" for more information.





(i) Note

The temporary directory must reside in the same directory path on each node in the cluster.

You must have 1 GB of disk space available in the TEMP directory. If you do not have sufficient space, then first delete all unnecessary files. If the temporary disk space is still less than the required amount, then increase the partition size of the disk or set the TEMP environment variable to point to a different hard drive. Ensure the environment variables TEMP and TMP both point to the location of the TEMP directory, for example:

TEMP=C:\WINDOWS\TEMP TMP=C:\WINDOWS\TEMP

- 1. From the Control Panel, select System.
- Select Advanced System Settings.
- In the System Properties windows, select the Advanced tab, then click Environment Variables.
- Modify the value of the TEMP environment variable in the user variables list.

General Server Requirements

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

- Select servers with the same instruction set architecture as cluster members.
- Ensure servers run the same operating system binary.
- Oracle Grid Infrastructure installations and Oracle Real Application Clusters (Oracle RAC) support servers with different hardware in the same cluster. Your cluster can have nodes with CPUs of different speeds or sizes, but Oracle recommends that you use nodes with the same hardware configuration.

Oracle recommends that if you configure clusters using different configuration, that you categorize cluster nodes into homogenous pools as part of your server categorization management policy.

Related Topics

Oracle Clusterware Administration and Deployment Guide

Server Minimum Hardware and Memory Requirements

Each system must meet certain minimum hardware and memory requirements.



Table 2-1 Minimum Hardware and Memory Requirements for Oracle Grid Infrastructure Installation

Hardware Component	Requirements
Memory (RAM)	Oracle Grid Infrastructure installations: At least 16 GB of physical RAM for Oracle Grid Infrastructure for a cluster installations, including installations where you plan to install Oracle RAC
Virtual memory (swap)	If your server has between 4 GB and 16 GB of RAM, then you should configure a paging file of at least the same size as the RAM. If your server has more than 16 GB of RAM, then the paging file should be 16 GB.
Video adapter	256 color and at least 1024 x 768 display resolution, so that OUI displays correctly while performing a system console-based installation
Processor	x64: Intel Extended Memory 64 Technology (EM64T) or AMD 64

(i) Note

32-bit systems are no longer supported for Oracle Grid Infrastructure and Oracle RAC.

Server Minimum Storage Requirements

Each system must meet certain minimum storage requirements.

1 GB of space in the %TEMP% directory.

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

- Delete unnecessary files from the %TEMP% directory to make available the space required.
- Extend the file system that contains the %TEMP% directory. If necessary, contact your system administrator for information about extending file systems.
- At least 12 GB of space for the Oracle Grid Infrastructure for a cluster home (Grid home).
 Oracle recommends that you allocate 100 GB to allow additional space for patches.
- At least 3.5 GB of space for the Oracle base of the Oracle Grid Infrastructure Installation user. The Oracle base includes Oracle Clusterware and Oracle ASM log files.
- If you intend to install Oracle Database, then allocate 7.5 GB of disk space for the Oracle home (the location for the Oracle Database software binaries).

If you plan to configure automated database backups for your Oracle Database, then you require additional space either in a file system or in an Oracle Automatic Storage Management disk group for the Fast Recovery Area.

Note

The base directory for Oracle Grid Infrastructure 12c and the base directory for Oracle RAC 12c must be different from the directories used by any Oracle RAC 11g Release 2 installations.



Disk Format Requirements

Oracle recommends that you install Oracle software, or binaries, on New Technology File System (NTFS) formatted drives or partitions.

Because it is difficult for OUI to estimate NTFS and file allocation table (FAT) disk sizes on Windows, the system requirements documented in this section are likely more accurate than the values reported on the OUI Summary screen.



(i) Note

Oracle Grid Infrastructure software is not supported on Network File System (NFS).

You cannot use NTFS formatted disks or partitions for Oracle Clusterware files or data files because they cannot be shared. Oracle Clusterware shared files and Oracle Database data files can be placed on unformatted basic disks or disk partitions, called raw partitions, managed by Oracle ASM.

Oracle Clusterware files and Oracle Database data files are stored on Oracle ASM, by default.

Configuring Time Synchronization for the Cluster

Oracle Clusterware requires the same time zone setting on all cluster nodes.

- **About Cluster Time Synchronization**
 - For the Microsoft Windows operating system, there are three different methods you can use to synchronize the time between the nodes of your cluster.
- **Understanding Network Time Requirements**
 - Oracle Clusterware 21c is automatically configured with Cluster Time Synchronization Service (CTSS).
- Configuring the Windows Time Service
 - The Windows Time service (W32Time) provides network clock synchronization on computers running Microsoft Windows.
- **Configuring Network Time Protocol**
 - The Network Time Protocol (NTP) is a client/server application.
- Configuring Cluster Time Synchronization Service
 - Cluster Time Synchronization Service (CTSS) is provided by Oracle to synchronize the time across the nodes in your cluster.

About Cluster Time Synchronization

For the Microsoft Windows operating system, there are three different methods you can use to synchronize the time between the nodes of your cluster.

During installation, the installation process picks up the time zone environment variable setting of the Oracle Installation user for Oracle Grid Infrastructure on the node where OUI runs. Then the installation process 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 three options for time synchronization between cluster nodes:

- Windows Time service
- An operating system configured network time protocol (NTP)
- 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 in observer mode. If neither NTP or the Windows Time service is found, then ctssd starts in active mode and synchronizes time among cluster members without contacting an external time server.

(i) 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.
- The IP address for an NTP server can be an IPv6 address.

Understanding Network Time Requirements

Oracle Clusterware 21c is automatically configured with Cluster Time Synchronization Service (CTSS).

CTSS provides automatic synchronization of the time settings on all cluster nodes. CTSS uses the optimal synchronization strategy for the type of cluster you deploy.

If you have an existing cluster synchronization service, such as network time protocol (NTP) or Windows Time Service, then CTSS starts in an observer mode. Otherwise, CTSS starts in an active mode to ensure that time is synchronized between cluster nodes. CTSS will not cause compatibility issues.

The CTSS module is installed as a part of Oracle Grid Infrastructure installation. CTSS daemons are started by the Oracle High Availability Services daemon (ohasd), and do not require a command-line interface.

Configuring the Windows Time Service

The Windows Time service (W32Time) provides network clock synchronization on computers running Microsoft Windows.

If you are using Windows Time service, and you prefer to continue using it instead of Cluster Time Synchronization Service, then you must modify the Windows Time service time jump settings, to allow the time to gradually match with the reference time. Restart the Windows Time service after you complete this task.

1. To configure Windows Time service, use the following command on each node:

C:\> W32tm /register

- To modify the Windows Time service to work in an Oracle RAC environment, perform the following steps:
 - a. Open the Registry Editor (regedit)



b. Locate the

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config key.

- c. Set the following Windows Time service parameters to these decimal values:
 - MaxPosPhaseCorrection to 600
 - MaxNegPhaseCorrection to 600
 - MaxAllowedPhaseOffset to 600

These parameter settings specify that small time adjustments are allowed when the time difference between the reference and cluster nodes is under 10 minutes.

① Note

You should configure the Windows Time service to meet the requirements of your environment, with assistance from Microsoft, if necessary. The recommended settings provided for the three parameters are the settings that Oracle recommends to allow time adjustments to happen through *slewing* (gradually adjusting the clock using small changes) rather than in large *steps* (setting the clock to a new time). Large time adjustments in a single step are not supported.

3. To put the changes into effect, use the following command:

C:\> W32tm /config /update

(i) See Also

For more information about using and configuring the Windows Time Service, see:

- Microsoft® Support article ID 816042: "How to configure an authoritative time server in Windows Server"
- Microsoft® Support article ID 939322: "Support boundary to configure the Windows Time service for high accuracy environments"
- NTP FAQ and HOW TO

Configuring Network Time Protocol

The Network Time Protocol (NTP) is a client/server application.

Each server must have NTP client software installed and configured to synchronize its clock to the network time server. The Windows Time service is not an exact implementation of the NTP, but it is based on the NTP specifications.

If you decide to use NTP instead of the Windows Time service, then, after you have installed the NTP client software on each cluster node, you must start the NTP service with the -x option to prevent time from being adjusted backward. Restart the network time protocol service after you complete this task:

- Use the registry editor to edit the value for the ntpd executable under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\NTP.
- 2. Add the -x option to the ImagePath key value, behind %INSTALLDIR%\ntpd.exe.



Restart the NTP service using the following commands:

net stop NTP net start NTP

Configuring Cluster Time Synchronization Service

Cluster Time Synchronization Service (CTSS) is provided by Oracle to synchronize the time across the nodes in your cluster.

When OUI discovers that neither the Windows Time or NTP services are active, the Cluster Time Synchronization Service is installed in active mode and synchronizes the time across the nodes. If the Windows Time service or NTP service is found on the server, then the Cluster Time Synchronization Service is started in observer mode, and no active time synchronization is performed by the Cluster Time Synchronization Service within the cluster.

To use Cluster Time Synchronization Service to provide synchronization service in the cluster, disable the Windows Time service and stop the NTP service. If you have an NTP service on your server but you cannot use the service to synchronize time with a time server, then you must deactivate and deinstall the NTP to use Cluster Time Synchronization Service.

 To confirm that the Cluster Time Synchronization Service is active after installation, enter the following command as the Oracle Grid Infrastructure installation owner:

crsctl check ctss

Configuring Operating Systems for Oracle Grid Infrastructure and Oracle RAC

Complete operating system configuration requirements and checks before you start the Oracle Grid Infrastructure installation.

- Reviewing Operating System and Software Upgrade Best Practices
 These topics provide general planning guidelines and platform-specific information about
- Reviewing Operating System Security Common Practices

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

upgrades and migration.

- Verifying Privileges for Copying Files in the Cluster During installation, OUI copies the software from the local node to the remote nodes in the cluster. The installation user must have privileges on the other nodes in the cluster to copy the files.
- Identifying Software Requirements
 Depending on the products that you intend to install, verify that the required operating system software is installed on each node of your cluster.
- Windows Firewall Feature on Windows Servers
 When installing Oracle Grid Infrastructure software or Oracle RAC software on Windows servers, it is mandatory to disable the Windows Firewall feature.
- Checking the Operating System Version
 To determine whether your computer is running a 64-bit (x64) Windows operating system, perform these steps.
- Checking Hardware and Software Certification on My Oracle Support
 The My Oracle Support website also provides compatible client and database releases, patches, and workaround information for bugs.
- 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.

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.
- Oracle ASM Upgrade Notifications
 Understand Oracle ASM upgrade options and restrictions.
- Rolling Upgrade Procedure Notifications
 Be aware of this information regarding rolling upgrades.



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

- <u>Upgrading Oracle Grid Infrastructure</u>
 Oracle Grid Infrastructure upgrade consists of upgrade of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM).
- Oracle Database Upgrade Guide

Oracle ASM Upgrade Notifications

Understand Oracle ASM upgrade options and restrictions.

- You can upgrade Oracle Automatic Storage Management (Oracle ASM) without shutting down an Oracle RAC database by performing a rolling upgrade either of individual nodes, or of a set of nodes in the cluster. However, if you have a single-instance database on a cluster that uses Oracle ASM, then you must shut down the single-instance database before upgrading.
- Oracle ASM is installed with Oracle Clusterware in the Oracle Grid Infrastructure home (Grid home).
- Two nodes of different releases cannot run in the cluster. When upgrading from Oracle
 Grid Infrastructure 12c release 2 (12.2), Oracle Grid Infrastructure 18c, or Oracle Grid
 Infrastructure 19c to a later release, if there is an outage during the rolling upgrade, then
 when you restart the upgrade, ensure that you start the earlier release of Oracle Grid
 Infrastructure and bring the Oracle ASM cluster back in the rolling migration mode.

Rolling Upgrade Procedure Notifications

Be aware of this information regarding rolling upgrades.

 During rolling upgrades of the operating system, Oracle supports using different operating system binaries when both versions of the operating system are certified with the Oracle Database release you are using.



- Using mixed operating system versions is supported during upgrade only.
 - Be aware that mixed operating systems are supported only for the duration of an upgrade, over the period of a few hours.
- Oracle Clusterware does not support nodes that have processors with different instruction set architectures (ISAs) in the same cluster. Each node must be binary compatible with the other nodes in the cluster.

For example, you cannot have one node using an Intel 64 processor and another node using an IA-64 (Itanium) processor in the same cluster. You could have one node using an Intel 64 processor and another node using an AMD64 processor in the same cluster because the processors use the same x86-64 ISA and run the same binary release of Oracle software.

(i) Note

Your cluster can have nodes with processors of different manufacturers, speeds, or sizes, but this is not recommended.

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.

Verifying Privileges for Copying Files in the Cluster

During installation, OUI copies the software from the local node to the remote nodes in the cluster. The installation user must have privileges on the other nodes in the cluster to copy the files.

1. Verify that you have Administrator privileges on the other nodes in the cluster by running the following command on each node, where *nodename* is the name of the remote node:

net use \\nodename\C\$

2. After installation, if your system does not use the net share shown in the above example, then you can remove the unused net share using the following command:

net use \\nodename\C\$ /delete

Identifying Software Requirements

Depending on the products that you intend to install, verify that the required operating system software is installed on each node of your cluster.

Requirements listed in this document are current as of the date listed on the title page. To obtain the most current information about operating system requirements, see the online version in the Oracle Help Center at the following URL:

https://docs.oracle.com/en/database/



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



(i) 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.

Table 3-1 Oracle Grid Software Requirements for Windows Systems

Requirement	Value
System Architecture	Processor: AMD64, or Intel Extended Memory (EM64T) Note: Oracle provides only x64 releases of Oracle Database with Oracle RAC for Windows.
	The x64 release of Oracle RAC runs on the x64 version of Windows on AMD64 and EM64T hardware.
Operating System	 Oracle Grid Infrastructure and Oracle RAC for x64 Windows: Windows Server 2022 x64 - Standard, Datacenter, and Essentials editions (Supported starting Oracle Database 21c Release Update 21.4 or later) Windows Server 2019 x64 - Standard, Datacenter, and Essentials editions Windows Server 2016 x64 - Standard, Datacenter, and Essentials editions Windows Server 2012 R2 x64 - Standard, Datacenter, Essentials, and Foundation editions Note: Windows Multilingual User Interface Pack is supported. The Server Core option is not supported.
Compiler and SDK	The following components are supported with the compilers based on Microsoft Visual C++ 2019 Update 7 (Version 16.7.8) or later and Intel 19.0 C Update 8 (Version 19.0.8), and Microsoft Visual C++ 2019 Update 7 (Version 16.7.8) or later SDK: Oracle Call Interface (OCI) External callouts Pro*C/C++ Oracle XML Developer's Kit (XDK) Oracle C++ Call Interface (OCCI) Pro*COBOL supports: Micro Focus Visual COBOL (Version 6)

Related Topics

Oracle Database Administrator's Reference for Microsoft Windows

Windows Firewall Feature on Windows Servers

When installing Oracle Grid Infrastructure software or Oracle RAC software on Windows servers, it is mandatory to disable the Windows Firewall feature.



If the Windows Firewall is enabled, then remote copy and configuration assistants such as virtual IP configuration assistant (VIPCA), Network Configuration Assistant (NETCA), and Oracle Database Configuration Assistant (DBCA) will fail during Oracle RAC installation. Thus, you must disable the firewall on all the nodes of a cluster before performing an Oracle RAC installation.

Note

The Windows Firewall should never be enabled on a NIC that is used as a cluster interconnect (private network interface) or for accessing an Oracle ASM network.

After the installation is successful, you can enable the Windows Firewall for the public connections. However, to ensure correct operation of the Oracle software, you must add certain executables and ports to the Firewall exception list on all the nodes of a cluster.

Additionally, the Windows Firewall must be disabled on all the nodes in the cluster before performing any clusterwide configuration changes, such as:

- Adding a node
- Deleting a node
- Upgrading to patch release
- Applying a patch bundle or an emergency patch

If you do not disable the Windows Firewall before performing these actions, then the changes might not be propagated correctly to all the nodes of the cluster.

Related Topics

Configuring Exceptions for the Windows Firewall

Checking the Operating System Version

To determine whether your computer is running a 64-bit (x64) Windows operating system, perform these steps.

- Right-click **My Computer** and select **Properties**.
- On the **General** tab, under the heading of System, view the displayed text.

You will see text similar to "64-bit Operating System" if you have the x64 version of the operating system installed.

Checking Hardware and Software Certification on My Oracle Support

The My Oracle Support website also provides compatible client and database releases, patches, and workaround information for bugs.

The hardware and software requirements included in this installation guide were current at the time this guide was 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.



View certification information on both Oracle Technology Network (OTN) and My Oracle Support.

The OTN certification page can be found on the following website: https://www.oracle.com/ database/real-application-clusters/

View My Oracle Support for guidance about supported hardware options that can assist you with your purchasing decisions and installation planning.

The My Oracle Support certifications page contains more detailed information about certified hardware and has information specific to each release and platform. My Oracle Support is available at the following URL:

https://support.oracle.com/

You must register online before using My Oracle Support. Use the steps described in the support document "Locate Oracle Database Server Certification Information for Microsoft Windows Platforms (Doc ID 1062972.1)" to locate the certification information for your Windows operating system.



(i) Note

Contact your Oracle sales representative if you do not have a My Oracle Support account.

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

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.

(i) Note

For the most up-to-date information about supported network protocols and hardware for Oracle RAC installations, refer to the Certify pages on the My Oracle Support website. See "Checking Hardware and Software Certification on My Oracle Support" for instructions.

- 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.
- <u>Understanding Network Addresses</u>
 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 Requirements

 Review these requirements to ensure that you have the minimum network hardware technology for Oracle Grid Infrastructure clusters.
- <u>Network Configuration Tasks for Windows Server Deployments</u>
 Microsoft Windows Server has many unique networking features. Some of these features require special configuration to enable Oracle software to run correctly on Windows Server.
- Oracle Grid Infrastructure IP Name and Address Requirements
 The Oracle Grid Naming Service (GNS) is used with large clusters to ease network administration cost.
- Intended Use of Network Adapters
 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.
- Broadcast Requirements for Networks Used by Oracle Grid Infrastructure
 Broadcast communications address resolution protocol (ARP) and User Datagram Protocol (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
 On each cluster member node the Oracle multicast DNS (mDNS) daemon uses multicasting on all network interfaces to communicate with other nodes in the cluster.
- Configuring Multiple ASM Interconnects on Microsoft Windows Platforms
 When using multiple network interface cards for the Oracle Automatic Storage
 Management (Oracle ASM) interconnect, you must enable the weakhostsend network
 parameter.



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 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. 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.





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

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

Related Topics

Oracle Clusterware Administration and Deployment Guide

About the SCAN

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

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

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

The SCAN resolves to multiple IP addresses reflecting multiple listeners in the cluster that handle public client connections. When a client submits a request, the SCAN listener listening on a SCAN IP address and the SCAN port is made available to a client. Because all services on the cluster are registered with the SCAN listener, the SCAN listener replies with the address of the local listener on the least-loaded node where the service is currently being offered. Finally, the client establishes connection to the service through the listener on the node where



service is offered. All of these actions take place transparently to the client without any explicit configuration required in the client.

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

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

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

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

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

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

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

Related Topics

Oracle Real Application Clusters Administration and Deployment Guide

Network Interface Hardware Requirements

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

- <u>Network Requirements for Each Node</u>
 Verify that servers where you install Oracle Grid Infrastructure meet the minimum network requirements for installation.
- <u>Network Requirements for the Private Network</u>
 The following is a list of requirements for the private network configuration:



- Network Requirements for the Public Network
 - The following is a list of requirements for the public network configuration:
- **IPv6 Protocol Support for Windows**
 - 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.
- Using Multiple Public Network Adapters
 - You can configure multiple network adapters for the public network interface.
- **Network Interface Configuration Options for Performance**
 - 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 Requirements for Each Node

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

- The host name of each node must use only the characters a-z, A-Z, 0-9, and the dash or minus sign (-). Host names using underscores () are not supported.
- Each node must have at least two network adapters or network interface cards (NICs): one for the public network interface, and one for the private network interface, or the **interconnect**. Each network adapter has a network connection name.



Note

Do not use the names PUBLIC and PRIVATE (all caps) for the public or private (interconnect) network connection names.

- Network adapters must be at least 1 GbE, with 10 GbE recommended.
- If you plan to use Oracle ASM running in a different cluster for storage, then you must either have a third network adapter for accessing the ASM network, or use the same network adapter that is used for the private network interface.

Network Requirements for the Private Network

The following is a list of requirements for the private network configuration:

- The private network connection names must be different from the network connection names used for the public network.
- Each node's private interface for interconnects must be on the same subnet.
 - 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]
- The private network connection name cannot contain any multibyte language characters. The private network connection names are case-sensitive.
- Both IPv4 and IPv6 addresses are supported.
 - 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.

- If you use OUI to install Oracle Grid Infrastructure, then the private network connection names associated with the private network adapters must be the same on all nodes.
 - For example, if you have a two-node cluster, and PrivNIC is the private network connection name for node1, then PrivNIC must be the private network connection name for node2.
- For the private network, the network adapters must use high-speed network adapters and switches that support UDP, TCP/IP (minimum requirement is 1 Gigabit Ethernet, 10GbE recommended). Alternatively, use InfiniBand for the interconnect.



(i) Note

UDP or TCP is the interconnect protocol for Oracle Clusterware. You must use a switch for the interconnect. Oracle recommends that you use a dedicated switch.

Oracle does not support token-rings or crossover cables for the interconnect.

For the private network adapters, the endpoints of all designated network connection names must be completely reachable on the network. There should be no node that is not connected to every other node on the private network. You can test if an interconnect interface is reachable using ping.

Network Requirements for the Public Network

The following is a list of requirements for the public network configuration:

- The public network connection names must be different from the private network connection names.
- Public network connection names are case-sensitive.
- The public network connection name cannot contain any multibyte language characters.
- If you use OUI to install Oracle Grid Infrastructure, then the public network connection names associated with the public network adapters for each network must be the same on all nodes.
 - For example, if you have a two-node cluster, you cannot configure network adapters on node1 with NIC1 as the public network connection name and on node2 have NIC2 as the public network connection name. Public network connection names must be the same, so you must configure NIC1 as the public network connection name on both nodes.
- For the public network, each network adapter must use high-speed network adapters and switches that support user datagram protocol (UDP), transmission control protocol and internet protocol (TCP/IP). Minimum requirement is 1 Gigabit Ethernet, 10GbE recommended.

IPv6 Protocol Support for Windows

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.



Cluster member node interfaces can be configured to use IPv4, IPv6, or both types of Internet protocol addresses. However, be aware of the following:

• Configuring public VIPs: 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. For example, 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 a network interface to use either the IPv4 protocol, or the IPv6 protocol on a given 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 configuring one interface to support IPv4 addresses and the other to support IPv6
 addresses. 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. Run the srvctl config network command to determine if DHCP or stateless address autoconfiguration is being used.

① Note

Link-local and site-local IPv6 addresses as defined in RFC 1884 are not supported

(i) See Also

- RFC 2732 for information about IPv6 notational representation
- RFC 3513 for information about proper IPv6 addressing
- RFC 2462 for information about IPv6 Stateless Address Autoconfiguration protocol
- Oracle Database Net Services Administrator's Guide for more information about network communication and IP address protocol options



Using Multiple Public Network Adapters

You can configure multiple network adapters for the public network interface.

Oracle recommends that you do not identify multiple public network connection names during Oracle Grid Infrastructure installation.

- 1. Use a third-party technology for your platform to aggregate the multiple public network adapters before you start installation.
- 2. During installation, select the single network connection name for the combined network adapters as the public interface.

If you configure two network adapters as public network adapters in the cluster without using an aggregation technology, the failure of one public network adapter on a node does not result in automatic VIP failover to the other public network adapter.

Network Interface Configuration Options for Performance

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.

If you access Oracle ASM remotely, or a 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 data communications. Failing to provide three separate network interfaces in this case can cause performance and stability problems under heavy system loads.

Network Configuration Tasks for Windows Server Deployments

Microsoft Windows Server has many unique networking features. Some of these features require special configuration to enable Oracle software to run correctly on Windows Server.

- <u>Disabling Windows Media Sensing</u>
 Windows Media Sensing must be disabled for the private network adapters.
- Deconfigure DNS Registration for Public Network Adapter
 To prevent Windows Server from potentially registering the wrong IP addresses for the node in DNS after a server restart, you must deconfigure the "Register this connection's addresses in DNS" option for the public network adapters.
- Manually Configure Automatic Metric Values
 Automatic Metric feature automatically configures the metric for the local routes that are based on link speed. To prevent OUI from selecting the wrong network interface during installation, you must customize the metric values for the public and private network interfaces.
- Setting UDP and TCP Dynamic Port Range for Oracle RAC Installations
 For certain configurations of Oracle RAC in high load environments it is possible for the system to exhaust the available number of sockets. To avoid this problem, expand the dynamic port range for both UDP and TCP.



Disabling Windows Media Sensing

Windows Media Sensing must be disabled for the private network adapters.

To disable Windows Media Sensing for TCP/IP, you must set the value of the DisableDHCPMediaSense parameter to 1 on each node. Because you must modify the Windows registry to disable Media Sensing, you should first backup the registry and confirm that you can restore it, using the methods described in your Windows documentation.

- Backup the Windows registry.
- 2. Use Registry Editor to view the following key in the registry:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters

3. Add a new DWORD value to the Parameters subkey:

Value Name: DisableDHCPMediaSense Value: 1

- 4. Exit the Registry Editor and restart the computer.
- 5. Repeat steps 1 through 4 on each node in your cluster.

Deconfigure DNS Registration for Public Network Adapter

To prevent Windows Server from potentially registering the wrong IP addresses for the node in DNS after a server restart, you must deconfigure the "Register this connection's addresses in DNS" option for the public network adapters.

Due to a change in functionality in the Windows Server 2008 operating system, the DNS client service registers all of the network connections of a computer in DNS. In earlier versions of Windows Server, the DNS client service registered only the primary, or first, network adapter IP address in DNS.

- Start the Windows Server Manager application.
- 2. Select View Network Connections.
- 3. Right-click the network adapter that provides the Public network interface and select **Properties**.
- 4. Select the Networking tab, and then select Internet Protocol Version 4 (TCP/IPv4).



If you configure this setting in IPv4, then Windows automatically configures the same setting for IPv6

- 5. Click Properties.
- 6. On the General tab, click Advanced.
- Select the DNS tab.
- 8. Deselect Register this connection's addresses in DNS.



Related Topics

DNS Best Practices

Manually Configure Automatic Metric Values

Automatic Metric feature automatically configures the metric for the local routes that are based on link speed. To prevent OUI from selecting the wrong network interface during installation, you must customize the metric values for the public and private network interfaces.

The Automatic Metric feature is enabled by default, and it can also be manually configured to assign a specific metric. The public and private network interface for IPv4 use the Automatic Metric feature of Windows. When the Automatic Metric feature is enabled and using the default values, it can sometimes cause OUI to select the private network interface as the default public host name for the server when installing Oracle Grid Infrastructure.

- 1. In Control Panel, double-click **Network Connections**.
- 2. Right-click a network interface, and then click **Properties**.
- 3. Click Internet Protocol (TCP/IP), and then click Properties.
- On the General tab, click Advanced.
- 5. To specify a metric, on the IP Settings tab, click to clear the Automatic metric check box.
- In the Interface Metric field, set the public network interface metric to a lower value than the private network interface. This assigns a higher priority to public network in the binding order.

For example, if you set the public network interface metric to 100 and the private network interface metric to 300, then the public network is assigned a higher priority than the private network.

Setting UDP and TCP Dynamic Port Range for Oracle RAC Installations

For certain configurations of Oracle RAC in high load environments it is possible for the system to exhaust the available number of sockets. To avoid this problem, expand the dynamic port range for both UDP and TCP.

- 1. Open a command line window as an Administrator user.
- Run the following commands to set the dynamic port range:

```
netsh int ipv4 set dynamicport udp start=9000 num=56000 netsh int ipv4 set dynamicport tcp start=9000 num=56000
```

3. Run the following commands to verify that the dynamic port range was set:

```
netsh int ipv4 show dynamicport udp
netsh int ipv4 show dynamicport tcp
```

For IPv6 network, replace IPv4 with IPv6 in the above examples.

Oracle Grid Infrastructure IP Name and Address Requirements

The Oracle Grid Naming Service (GNS) is used with large clusters to ease network administration cost.



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.

- 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.
- <u>Cluster Name and SCAN Requirements</u>
 Review this information before you select the cluster name and SCAN.
- IP Name and Address Requirements For Grid Naming Service (GNS)
 The network administration must configure the domain name server (DNS) to delegate resolution requests for cluster names (any names in the subdomain delegated to the cluster) to the 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 Address Requirements for Manual Configuration
 If you do not enable GNS, then the public and VIP addresses for each node must be static
 IP addresses. Public, VIP and SCAN addresses must be on the same subnet.
- Confirming the DNS Configuration for SCAN
 You can use the nslookup command to confirm that the DNS is correctly associating the SCAN with the addresses.
- Grid Naming Service Cluster Configuration Example
 Review this example to understand Grid Naming Service configuration.
- <u>Domain Delegation to Grid Naming Service</u>
 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.
- Manual IP Address Configuration Example
 If you choose not to use GNS, then before installation you must configure public, virtual, and private IP 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.

You can configure IP addresses for Oracle Grid Infrastructure and Oracle RAC 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 to be resolved by GNS (standard or multicluster). Requirements for GNS are different depending on whether you choose to configure GNS with zone delegation (resolution of a domain delegated to GNS), or without zone delegation (a GNS virtual IP address without domain delegation).
- 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 DHCP addresses to provide 1 IP address for each node, and 3 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), a single GNS instance can be used by multiple clusters. 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 the Oracle Clusterware managed VIPs. In addition, domain
 name system (DNS)-based static name resolution is used for each node, or host files for
 both the clusters and clients have to be updated, and SCAN functionality is limited.
 Selecting this option requires that you request network administration updates when you
 modify the cluster.

(i) Note

- Oracle recommends that you use a static host name for all non-VIP server node public host names.
- Public IP addresses and virtual IP addresses must be in the same subnet.
- Oracle only supports DHCP-assigned networks for the default network, not for any subsequent networks.

For clusters using single interfaces for private networks, each node's private interface for interconnects must be on the same subnet, and that 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].

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.

(i) 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)

The network administration must configure the domain name server (DNS) to delegate resolution requests for cluster names (any names in the subdomain delegated to the cluster) to the GNS.

If you enable Grid Naming Service (GNS), then name resolution requests to the cluster are delegated to the GNS, which listens on the GNS VIP address. 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.

(i) Note

You cannot use GNS with another multicast DNS. To use GNS, disable any third-party mDNS daemons on your system.

Related Topics

Configuring DNS for Domain Delegation to Grid Naming Service

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. Multicluster 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 | secondary}

For example:

C:\> srvctl export gns -clientdata C:\Users\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.



- 3. During installation, you are prompted to provide a path to that file.
- 4. After you have completed the GNS Client cluster installation, you must run the following command on one of the GNS Server cluster members to start GNS service, where path_to_file is the name and path location of the GNS Client data file:

srvctl add gns -clientdata path_to_file

For example:

C:\> srvctl add gns -clientdata C:\Users\grid\gns_client_data

Related Topics

Oracle Clusterware Administration and Deployment Guide

IP Address Requirements for Manual Configuration

If you do not enable GNS, then the public and VIP addresses for each node must be static IP addresses. Public, VIP and SCAN addresses must be on the same subnet.

IP addresses on the subnet you identify as private are assigned as private IP addresses for cluster member nodes. Oracle Clusterware manages private IP addresses in the private subnet. You do not have to configure these addresses manually in a hosts file.

The cluster must have the following addresses configured:

- A public IP address for each node configured before installation, and resolvable to that node before installation.
- A VIP address for each node configured before installation, but not currently in use (or not resolvable to that node before installation).
- 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 by the DNS to the requestor. These addresses must be configured before installation in the DNS to resolve to addresses that are not currently in use. The SCAN name must meet the requirements specified in "Cluster Name and SCAN Requirements".
- A private IP address for each node configured before installation, but on a separate, private network, with its own subnet. The IP address should not be resolvable except by other cluster member nodes.
- A set of one or more networks over which Oracle ASM serves its clients. The ASM network
 does not have to be a physical network; it can be a virtual network. The ASM network must
 use either a third NIC, or share a private network adapter. The NIC can be a virtual NIC.

① Note

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 you will only be able to resolve to one IP address and you will have only one SCAN address.

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.



Related Topics

Understanding Network Addresses

Confirming the DNS Configuration for SCAN

You can use the nslookup command to confirm that the DNS is correctly associating the SCAN with the addresses.

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

C:\> nslookup mycluster-scan Server: dns3.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.

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	Туре	Address	Address Assigned By	Resolved By
GNS VIP	None	Selected by Oracle Clusterware	mycluster- gns- vip.example.c om	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.myclust er.cluster01.e xample.com	virtual	192.0.2.201	DHCP	GNS
SCAN VIP 2	none	Selected by Oracle Clusterware	mycluster- scan.myclust er.cluster01.e xample.com	virtual	192.0.2.202	DHCP	GNS
SCAN VIP 3	none	Selected by Oracle Clusterware	mycluster- scan.myclust er.cluster01.e xample.com	virtual	192.0.2.203	DHCP	GNS

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 Domain Delegation to Grid Naming Service
 You must configure the DNS to send GNS name resolution requests using DNS forwarders.



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 racgns.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 Domain Delegation to Grid Naming Service

You must configure the DNS to send GNS name resolution requests using DNS forwarders.

If the DNS server is running on a Windows server that you administer, then the following steps must be performed to configure DNS:

Click Start, then select Programs. Select Administrative Tools and then click DNS manager.

The DNS server configuration wizard starts automatically.

Use the wizard to create an entry for the GNS virtual IP address, where the address is a valid DNS name.

For example, if the cluster name is mycluster, and the domain name is example.com, and the IP address is 192.0.2.1, you could create an entry similar to the following:

mycluster-gns-vip.example.com: 192.0.2.1

The address you provide must be routable.



(i) Note

The domain name may not contain underscores. Windows may allow the use of underscore characters, but this practice violates the Internet Engineering Task Force RFC 952 standard and is not supported by Oracle.

- To configure DNS forwarders, click Start, select Administrative Tools, and then select DNS.
- 4. Right-click **ServerName**, where **ServerName** is the name of the server, and then click the Forwarders tab.
- 5. Click **New**, then type the name of the DNS domain for which you want to forward queries in the DNS domain box, for example, clusterdomain.example.com. Click OK.
- In the selected domain's forwarder IP address box, type the GNS VIP address, and then click Add.



7. Click OK to exit.

(i) Note

Experienced DNS administrators may want to create a reverse lookup zone to enable resolution of reverse lookups. A **reverse lookup** resolves an IP address to a host name with a Pointer Resource (PTR) record. If you have reverse DNS zones configured, then you can automatically create associated reverse records when you create your original forward record.

(i) See Also

Oracle Grid Infrastructure Installation Guide for your platform or your operating system documentation for more information about configuring DNS for domain delegation to GNS if the DNS server is running on a different operating system.

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	Туре	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



Table 4-2 ((Cont.)	Manual	Network	Config	uration	Example

Identity	Home Node	Host Node	Given Name	Туре	Address	Address Assigned By	Resolved By
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.



(i) Note

All host names must conform to the RFC-952 standard, which permits alphanumeric characters, but does not allow underscores (" ").

Intended Use of Network Adapters

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 performs only one of the following roles:

- **Public**
- Private
- **ASM**
- **ASM** and Private
- Do Not Use

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.

For network adapters that you plan to use for other purposes, for example, an adapter dedicated to a non-Oracle network file system, you must identify those network adapters as "do not use" adapters so that Oracle Clusterware ignores them.

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.

Oracle Flex ASM can use either the same private networks as Oracle Clusterware, or use its own dedicated private networks. If you plan to have other clusters (Oracle ASM client clusters) access the storage in an Oracle Flex ASM cluster, then you should configure a separate ASM network as a private network that connects the client clusters to the Oracle Flex ASM cluster.

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.

Broadcast Requirements for Networks Used by Oracle Grid Infrastructure

Broadcast communications address resolution protocol (ARP) and User Datagram Protocol (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 ARP, which is necessary for VIP failover. Do not configure NOARP.

Multicast Requirements for Networks Used by Oracle Grid Infrastructure

On each cluster member node the Oracle multicast DNS (mDNS) daemon uses multicasting on all network interfaces to communicate with other nodes in the cluster.

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

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

You do not need to enable multicast communications across routers.

Configuring Multiple ASM Interconnects on Microsoft Windows Platforms

When using multiple network interface cards for the Oracle Automatic Storage Management (Oracle ASM) interconnect, you must enable the weakhostsend network parameter.

Microsoft Windows versions prior to Windows Vista use the **weak host send and receive model**. The TCP/IP stack in Windows Vista, Windows Server 2008, and later Windows operating systems, supports the **strong host send and receive model** for both IPv4 and IPv6 protocols by default. Oracle RAC nodes that use multiple network interface cards (NICs) for the ASM interconnect must enable the weak host send by specifically setting the weakhostsend parameter for all the private subnets.

If you do not enable the <code>weakhostsend</code> parameter, you might experience connection issues on the ASM interconnects because the interconnect packets are being blocked or discarded. It is not



considered unsafe to enable weakhostsend because the interconnect is on a private and isolated network. Enabling the weakhostsend parameter allows all private NICs to send packets to multiple private subnets.

- Open a CMD Prompt as Administrator window.
- Use the following commands to configure the send behavior for both IPv4 and IPv6, on a
 per-interface basis, where interface1 and interface2 represent the names or interface
 indexes assigned to the NIC adapters on the Oracle RAC nodes.

netsh interface [ipv4 | ipv6] set interface *interface1* weakhostsend=enabled netsh interface [ipv4 | ipv6] set interface *interface2* weakhostsend=enabled

You can obtain the interface index of a NIC from the output of the command:

netsh interface [ipv4 | ipv6] show interface

3. Repeat the commands in Step 2 for all the private NICs, on each node in the cluster.

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

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

You must configure certain users, groups, and environment settings used during Oracle Grid Infrastructure for a Cluster and Oracle Real Application Clusters installations.

- <u>Creating Installation Groups and Users for Oracle Grid Infrastructure and Oracle RAC</u>
 To install Oracle Grid Infrastructure and Oracle RAC, you must have an installation user and optionally an Oracle Home user for Oracle RAC.
- <u>Standard Administration and Job Role Separation User Groups</u>
 Oracle Grid Infrastructure uses various operating system groups.
- Configuring User Accounts
 When installing Oracle Grid Infrastructure for a cluster, you run the installer software as an Administrator user.
- <u>Creating Oracle Software Directories</u>
 During installation, you are prompted to provide a path to a home directory to store Oracle
 Grid Infrastructure software.

Related Topics

• <u>User Environment Configuration Checklist for Oracle Grid Infrastructure</u> Review the following environment checklist for all installations.

Creating Installation Groups and Users for Oracle Grid Infrastructure and Oracle RAC

To install Oracle Grid Infrastructure and Oracle RAC, you must have an installation user and optionally an Oracle Home user for Oracle RAC.



During Oracle Grid Infrastructure installation, both Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) are installed. You no longer can have separate Oracle Clusterware installation owners and Oracle ASM installation owners.

- About the Oracle Installation User
 The Oracle Installation User can be either a local user or a domain user.
- About the Oracle Home User for the Oracle RAC Installation
 You can specify an optional Oracle Home user for Oracle RAC installation and configuration.
- When to Create an Oracle Home User
 You must create an Oracle Home User in certain circumstances.



- Oracle Home User Configurations for Oracle Installations When the Oracle software installation completes, you will have one of the following configurations:
- Understanding the Oracle Inventory Directory and the Oracle Inventory Group You must have a group whose members are given access to write to the Oracle Inventory directory, which is the central inventory record of all Oracle software installations on a

About the Oracle Installation User

The Oracle Installation User can be either a local user or a domain user.

To install the Oracle Grid Infrastructure or Oracle Database software, you must use either a local or domain user. In either case, the Oracle Installation User must be an explicit member of the Administrators group on all nodes of the cluster.

If you use a local user account for installing Oracle Grid Infrastructure or Oracle Real Application Clusters (Oracle RAC), then:

- The user account must exist on all nodes in the cluster.
- The user name and password must be the same on all nodes.
- OUI displays a warning message.

If you use a domain user account for installing Oracle Grid Infrastructure or Oracle Real Application Clusters (Oracle RAC), then:

- The domain user must be explicitly declared as a member of the local Administrators group on each node in the cluster. It is not sufficient if the domain user has inherited membership from another group.
- The user performing the installation must be in the same domain on each node. For example, you cannot have the DBADMIN\dba1 user on the first node and the RACDBA\dba1 user on the second node.
- A local user of the same name cannot exist on any of the nodes. For example if you use RACDBA\dba1 as the installation user, none of the nodes can have a local NODE1\dba1 user account.

About the Oracle Home User for the Oracle RAC Installation

You can specify an optional Oracle Home user for Oracle RAC installation and configuration.



(i) Note

Starting with Oracle Database 21c, you cannot specify an Oracle Home user during Oracle Grid Infrastructure installation.

During installation of Oracle RAC, you can either use a Windows built-in account or specify an optional, non-Administrator user that is a Windows domain user or a Windows Group Managed Service Account (gMSA) to be the **Oracle Home User** associated with the Oracle RAC home.

If a Windows domain user account or a gMSA is chosen, then it should be an existing domain user account with no administration privileges.



For Oracle RAC installations, Oracle recommends that you use a Windows domain user or a qMSA (instead of Windows built-in account) as the Oracle Home User for enhanced security.

The services created for the Oracle RAC software run using the privileges of the Oracle Home User for Oracle RAC, or the Local System built-in Windows account if you did not specify an Oracle Home User during installation. Oracle Universal Installer (OUI) creates multiple operating system groups, such as the ORA DBA group, on all nodes. The user performing the installation is automatically added to those groups necessary for proper database administration.

For an administrator-managed database, you have the option of storing Oracle Home User password in a secure wallet (stored in Oracle Cluster Registry). Use the following CRSCTL command to create this secure wallet for storing the Windows operating system user name and password:

crsctl add wallet -osuser -passwd

If the wallet (stored in Oracle Cluster Registry) exists, then Oracle administration tools automatically use the password from the wallet without prompting the administrator to enter the password of Oracle Home User for performing administrative operations.

A policy-managed database mandates the storage of Oracle Home User password in the wallet (stored in Oracle Cluster Registry). When a policy-managed database is created, DBCA automatically creates the wallet, if one does not exist.



(i) Note

If you choose to use an Oracle Home User for your Oracle RAC installation, then the Windows User Account you specify must be a domain user or a gMSA.

Related Topics

Oracle Database Administrator's Reference for Microsoft Windows

When to Create an Oracle Home User

You must create an Oracle Home User in certain circumstances.

If an Oracle Home User exists, but you want to use a different operating system user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation.

- Restrictions and Guidelines for Oracle Home Users Review the following restrictions and guidelines for Oracle Home Users for Oracle software installations.
- Determining if an Oracle Home User Exists You must decide to use an existing user, or create a new user.
- Creating an Oracle Home User Use the Manage User Accounts window to create a new user.
- Using an Existing Oracle Software Owner User If the user you have decided to use as an Oracle Home user exists, then you can use this user as the Oracle Home user for a different installation.



Restrictions and Guidelines for Oracle Home Users

Review the following restrictions and guidelines for Oracle Home Users for Oracle software installations.

- If you plan to install Oracle Database or Oracle RAC, then Oracle recommends that you create a separate Oracle Home user for the Oracle Database installation. If you use one Oracle Home user, then when you want to perform administration tasks, you must select the utilities from the Oracle home for the instance you want to administer, or change the default %ORACLE_HOME% value to the location of the Oracle Home from which the instance runs. For Oracle ASM instances, you must use the Oracle Grid Infrastructure home and for database instance use the Oracle Database home.
- If you try to administer an Oracle home or Grid home instance using sqlplus, srvctl, Isnrctl, 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, SRVTCL ignores %ORACLE_HOME%, and the Oracle home environment variable does not affect SRVCTL commands. In all other cases, you must start the utilities from the Oracle home of the instance that you want to administer.

If you need to set the user environment to use a specific Oracle home, then use Oracle Universal Installer. On the landing page, click **Installed Products**. In the Inventory window, click the **Environment** tab. Select the Oracle Home you want to use, and deselect the other Oracle homes, then click **Apply**. You can then exit Oracle Universal Installer. When you use Oracle Universal Installer to set the Oracle Home, it updates the ORACLE_HOME environment variable and updates the PATH variable.

Determining if an Oracle Home User Exists

You must decide to use an existing user, or create a new user.

- 1. Open the Control Panel window.
- Select User Accounts.
- 3. Select Manage User Accounts.
- Scroll through the list of names until you find the ones you are looking for.If the names do not appear in the list, then the user has not yet been created.

Creating an Oracle Home User

Use the Manage User Accounts window to create a new user.

Perform the following steps to create an Oracle Home user on the domain server with domain administrator privileges.

The user must not be a member of the Administrators group. If you are creating an Oracle Home User for an Oracle RAC installation, then the user must be a Windows domain user, and the user must be a member of the same domain on each node in the cluster.

- Open the Control Panel window
- 2. Select User Accounts.
- Select Manage User Accounts.
- 4. Create the user using the interface.



Related Topics

Oracle Database Administrator's Reference for Microsoft Windows

Using an Existing Oracle Software Owner User

If the user you have decided to use as an Oracle Home user exists, then you can use this user as the Oracle Home user for a different installation.

Oracle does not support changing the ownership of an existing Oracle Database home from one Oracle Home user to a different user.

During the software installation, specify the existing user for the Oracle Home user. Oracle Universal Installer (OUI) creates the appropriate group memberships.

Oracle Home User Configurations for Oracle Installations

When the Oracle software installation completes, you will have one of the following configurations:

Installation Type	Oracle Home user configuration
Oracle RAC with specified Oracle Home user	The Oracle Home user owns all the services run by the Oracle Database software.
Oracle RAC with Built-in Oracle Home user	The services run under the built-in LocalSystem account.



Note

You cannot change the Oracle Home User after installation to a different Oracle Home user. Only out-of-place upgrade or move allows the Oracle Home user to be changed to or from the built-in Windows account.

Understanding the Oracle Inventory Directory and the Oracle Inventory Group

You must have a group whose members are given access to write to the Oracle Inventory directory, which is the central inventory record of all Oracle software installations on a server.

When you install Oracle software on the system for the first time, Oracle Universal Installer (OUI) creates the directories for the Oracle central inventory. OUI also creates the Oracle Inventory group, ORA_INSTALL. The ORA_INSTALL group contains all the Oracle Home users for all Oracle homes on the server. The location of the Oracle central inventory on Windows is always %SYSTEM DRIVE%\Program Files\Oracle\Inventory.

Whether you are performing the first installation of Oracle software on this server, or are performing an installation of additional Oracle software on the server, you do not need to create the Oracle central inventory or the ORA_INSTALL group. You cannot change the name of the Oracle Inventory group - it is always ORA_INSTALL.

Members of the Oracle Inventory group have write privileges to the Oracle central inventory 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. All Oracle software install users must be members of the Oracle Inventory group. Members of this group can talk to Cluster Synchronization Service (CSS).

(i) Note

If Oracle software is already installed on the system, then, when you install new Oracle software, the existing Oracle Inventory group is used instead of creating a new Inventory group.

Standard Administration and Job Role Separation User Groups

Oracle Grid Infrastructure uses various operating system groups.

These operating system groups are designated with the logical role of granting operating system group authentication for administration system privilege for Oracle Clusterware and Oracle ASM.

- About Job Role Separation Operating System Privileges Groups and Users
 Job role separation requires that you create different operating system groups for each set of system privileges that you grant through operating system authorization.
- Oracle Software Owner for Each Oracle Software Product
 Oracle recommends that you use appropriate operating system groups and users for all installations where you specify separate Oracle Home Users:
- <u>Standard Oracle Database Groups for Database Administrators</u>
 The Oracle Database supports multiple operating system groups to provide operating system authentication for database administration system privileges.
- Oracle ASM Groups for Job Role Separation
 The SYSASM, SYSOPER for ASM, and SYSDBA for ASM system privileges enables the separation of the Oracle ASM storage administration privileges from SYSDBA.
- Extended Oracle Database Administration Groups for Job Role Separation
 Oracle Database 12c Release 1 (12.1) and later releases provide an extended set of database groups to grant task-specific system privileges for database administration.
- Operating System Groups Created During Installation
 When you install either Oracle Grid Infrastructure or Oracle RAC, the user groups listed in the following table are created, if they do not already exist.
- Example of Using Role-Allocated Groups and Users
 You can use role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment.

About Job Role Separation Operating System Privileges Groups and Users

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 authentication for Oracle ASM system privileges for storage tier administration. This operating system authentication is separated from Oracle Database operating system authentication. In addition, the Oracle Grid Infrastructure Installation user provides operating system user authentication for modifications to Oracle Grid Infrastructure binaries.



With Oracle Database job role separation, each Oracle Database installation has separate operating system groups. The operating system groups provide authorization for system privileges on that Oracle Database, so multiple databases can be installed on the cluster without sharing operating system authentication for system privileges. In addition, each Oracle software installation is associated with an Oracle Installation user, to provide operating system user authorization for modifications to Oracle Database binaries.

Note

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.

During the Oracle Database installation, the installation creates the OSDBA, OSOPER, OSBACKUPDBA, OSDGDBA, OSKMDBA, and OSRACDBA groups and you can assign users to these groups. Members of these groups are granted operating system authentication for the set of database system privileges each group authorizes. Oracle recommends that you use different operating system groups for each set of system privileges.

(i) Note

This configuration is optional, to restrict user access to Oracle software by responsibility areas for different administrator users.

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

Oracle Software Owner for Each Oracle Software Product

Oracle recommends that you use appropriate operating system groups and users for all installations where you specify separate Oracle Home Users:

You can choose from the following operating system groups and users, as per your configuration:

Separate Oracle Installation users for each Oracle software product (typically, oracle, for the Oracle Database software, and grid for the Oracle Grid Infrastructure software.

An Oracle installation user must have Administrator privilege to install Oracle sofware on the system. This user owns the Oracle binaries of the Oracle Grid Infrastructure software, and you can also use this same user as the Oracle Installation user for the Oracle Database or Oracle RAC binaries.

The Oracle Installation user for Oracle Database software has full administrative privileges for Oracle instances and is added to the ORA DBA, ORA ASMDBA, ORA HOMENAME SYSBACKUP, ORA HOMENAME SYSDG, ORA HOMENAME SYSKM, and ORA HOMENAME SYSRAC groups. Oracle Home users are added to the ORA_HOMENAME_DBA group for the Oracle home created during



the installation. The ORA_OPER and ORA_HOMENAME_OPER groups are created, but no users are added to these groups during installation.

Related Topics

Oracle Database Security Guide

Standard Oracle Database Groups for Database Administrators

The Oracle Database supports multiple operating system groups to provide operating system authentication for database administration system privileges.

OSDBA group (ORA_DBA)

When you install Oracle Database, a special Windows local group called ORA_DBA is created (if it does not already exist from an earlier Oracle Database installation), and the Oracle Installation user is automatically added to this group. Members of the ORA_DBA group automatically receive the SYSDBA privilege. Membership in the ORA_DBA group allows a user to:

- Connect to Oracle Database instances without a password
- Perform database administration procedures such as starting and shutting down local databases
- Add additional Windows users to ORA_DBA, enabling them to have the SYSDBA privilege

Membership in the ORA_DBA group grants full access to all databases on the server.

OSDBA group for a particular Oracle home (ORA_HOMENAME_DBA)

This group is created the first time you install Oracle Database software into a new Oracle home. Membership in the ORA_HOMENAME_DBA group grants full access (SYSDBA privileges) for all databases that run from the specific Oracle home.

Belonging to either the ORA_DBA or ORA_HOMENAME_DBA group does not grant any special privileges for the user with respect to the Oracle ASM instance. Members of these groups will not be able to connect to the Oracle ASM instance.

OSOPER group for Oracle Database (ORA_OPER)

This group is created the first time you install Oracle Database software into a new Oracle home. This optional group identifies operating system user accounts that have database administrative privileges (the SYSOPER system privilege) for the database instances that run from any Oracle home. Assign users to this group if you want a separate group of operating system users to have a limited set of database administrative privileges for starting up and shutting down any Oracle database.

OSOPER group for a particular Oracle home (ORA_HOMENAME_OPER)

This group is created the first time you install Oracle Database software into a new Oracle home. This optional group identifies operating system user accounts that have database administrative privileges (the SYSOPER system privilege) for the database instances that run from a specific Oracle home. Assign users to this group if you want a separate group of operating system users to have a limited set of database administrative privileges for starting up and shutting down any Oracle database located in a specific Oracle home.



Oracle ASM Groups for Job Role Separation

The SYSASM, SYSOPER for ASM, and SYSDBA for ASM system privileges enables the separation of the Oracle ASM storage administration privileges from SYSDBA.

During installation, the following groups are created for Oracle ASM:

OSASM Group for Oracle ASM Administration (ORA_ASMADMIN)

Use this separate group to have separate administration privilege groups for Oracle ASM and Oracle Database administrators. Members of this group are granted the SYSASM system privilege to administer Oracle ASM. In Oracle documentation, the operating system group whose members are granted privileges is called the OSASM group. During installation, the Oracle Installation User for Oracle Grid Infrastructure and Oracle Database Service IDs are configured as members of this group. Membership in this group also grants database access to the ASM disks.

Members of the OSASM group can use SQL to connect to an Oracle ASM instance as SYSASM using operating system authentication. The SYSASM system privilege permits mounting and dismounting disk groups, and other storage administration tasks. SYSASM system privileges do not grant access privileges on an Oracle Database instance.

OSDBA for ASM Database Administrator group (ORA_ASMDBA)

This group grants access for the database to connect to Oracle ASM. During installation, the Oracle Installation Users are configured as members of this group. After you create an Oracle Database, this groups contains the Oracle Home Users of those database homes.

OSOPER for ASM Group for ASM Operators (ORA_ASMOPER)

This is an optional group. Use this group if you want a separate group of operating system users to have a limited set of Oracle ASM instance administrative privileges (the SYSOPER for ASM system 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 system privilege.

To use the Oracle ASM Operator group to create an Oracle ASM administrator with fewer privileges than those granted by the SYSASM system privilege you must assign the user to this group after installation.

Extended Oracle Database Administration Groups for Job Role Separation

Oracle Database 12c Release 1 (12.1) and later releases provide 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 ORA_DBA/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.

The installer automatically creates operating system groups whose members are granted these system privileges. The subset of OSDBA job role separation privileges and groups consist of the following:

OSBACKUPDBA group for Oracle Database (ORA HOMENAME SYSBACKUP)

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



OSDGDBA group for Oracle Data Guard (ORA HOMENAME SYSDG)

Assign users to 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.

OSKMDBA group for encryption key management (ORA_HOMENAME_SYSKM)

Assign users to 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.

 OSRACDBA group for Oracle Real Application Clusters Administration (typically, ORA_HOMENAME_SYSRAC)

Assign users to 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.

You cannot change the name of these operating system groups. These groups do not have any members after database creation, but an Administrator user can assign users to these groups after installation. Each operating system group identifies a group of operating system users that are granted the associated set of database privileges.

(i) See Also

- Oracle Database Administrator's Guide for more information about the OSDBA, OSASM, OSOPER, OSBACKUPDBA, OSDGDBA, OSKMDBA, and OSRACDBA groups, and the SYSDBA, SYSASM, SYSOPER, SYSBACKUP, SYSDG, SYSKM, and SYSRAC privileges
- The "Managing Administrative Privileges" section in Oracle Database Security Guide

Operating System Groups Created During Installation

When you install either Oracle Grid Infrastructure or Oracle RAC, the user groups listed in the following table are created, if they do not already exist.

Table 5-1 Operating System Groups Created During Installation

Operating System Group Names	Database Privileges	Description
ORA_ASMADMIN	SYSASM system	The OSASM group for the Oracle ASM instance.
	privileges for Oracle ASM administration	Using this group and the SYSASM system privileges enables the separation of SYSDBA database administration privileges from Oracle ASM storage administration privileges. Members of the OSASM group are authorized to connect using the SYSASM privilege and have full access to Oracle ASM, including administrative access to all disk groups that the Oracle ASM instance manages.



Table 5-1 (Cont.) Operating System Groups Created During Installation

Operating System Group Names	Database Privileges	Description
ORA_ASMDBA	SYSDBA system	The OSDBA group for the Oracle ASM instance.
	privileges on the Oracle ASM instance	This group grants access for the database to connect to Oracle ASM. During installation, the Oracle Installation Users are configured as members of this group. After you create an Oracle Database, this groups contains the Oracle Home Users of those database homes.
ORA_ASMOPER	SYSOPER for	The OSOPER group for the Oracle ASM instance.
	Oracle ASM system privileges	Members of this group are granted SYSOPER system privileges on the Oracle ASM instance, which permits a user to perform operations such as startup, shutdown, mount, dismount, and check disk group. This group has a subset of the privileges of the OSASM group.
		Similar to the ORA_HOMENAME_OPER group, this group does not have any members after installation, but you can manually add users to this group after the installation completes.
ORA_ <i>GRIDHM</i> _DBA	SYSDBA system privileges for the Oracle Grid Infrastructure	Members of this group are granted the SYSDBA system privileges for managing the Oracle Grid Infrastructure Management Repository database, where <i>GRIDHM</i> is the name of the Oracle Grid Infrastructure home.
	Management Repository database	The default home name is OraGrid12Home1, so the default group name is ORA_OraGrid12Home1_DBA.
ORA_GRIDHM_OPER	SYSOPER system privileges for the Oracle Grid Infrastructure	Members of this group are granted the SYSOPER system privileges for managing the Oracle Grid Infrastructure Management Repository database, where <i>GRIDHM</i> is the name of the Oracle Grid Infrastructure home.
	Management Repository database	If you use the default Grid home name of OraGrid12Home1,then the default operating system group name is ORA_OraGrid12Home1_OPER.
ORA_DBA	SYSDBA system	A special OSDBA group for the Windows operating system.
	privileges for all Oracle Database installations on the server	Members of this group are granted SYSDBA system privileges for all Oracle Databases installed on the server.
ORA_OPER	SYSOPER system	A special OSOPER group for the Windows operating system.
	privileges for all Oracle databases installed on the server	Members of this group are granted SYSOPER system privileges all Oracle Databases installed on the server. This group does not have any members after installation, but you can manually add users to this group after the installation completes.
ORA_ <i>HOMENAME</i> _DBA	SYSDBA system privileges for all	An OSDBA group for a specific Oracle Home with a name of HOMENAME.
	database instances that run from the Oracle home with the name HOMENAME	Members of this group can use operating system authentication to gain SYSDBA system privileges for any database that runs from the specific Oracle home. If you specified an Oracle Home User during installation, the user is added to this group during installation.



Table 5-1 (Cont.) Operating System Groups Created During Installation

Operating System Group Names	Database Privileges	Description		
ORA_ <i>HOMENAME</i> _OPER	SYSOPER system privileges for all	An OSDBA group for the Oracle Home with a name of HOMENAME.		
	database instances that run from the Oracle home with the name HOMENAME	Members of this group can use operating system authentication to gain SYSOPER system privileges for any database that runs from the specific Oracle home. This group does not have any members after installation, but you can manually add users to this group after the installation completes.		
ORA_ <i>HOMENAME</i> _SYSBACKUP	SYSBACKUP system privileges for	OSBACKUPDBA group for a specific Oracle Home with a name of <i>HOMENAME</i> .		
	all database instances that run from the Oracle home with a name of HOMENAME	Members of this group have privileges necessary for performing database backup and recovery tasks on all database instances that run from the specified Oracle Home directory.		
ORA_ <i>HOMENAME</i> _SYSDG	SYSDG system privileges for all	OSDGDBA group for a specific Oracle Home with a name of <i>HOMENAME</i> .		
	database instances that run from the Oracle home with a name of HOMENAME	Members of this group have privileges necessary for performing Data Guard administrative tasks on all database instances that run from the specified Oracle Home directory.		
ORA_ <i>HOMENAME</i> _SYSKM	SYSKM system privileges for all	OSKMDBA group for a specific Oracle Home with a name of <i>HOMENAME</i> .		
	database instances that run from the Oracle home with a name of HOMENAME.	Members of this group have privileges necessary for performing encryption key management tasks on all database instances that run from the specified Oracle Home directory.		
ORA_CRS_USERS	None	Members of this group have privileges necessary for file system permissions on the Grid Infrastructure Oracle Base directory.		
		When you configure a CRS wallet of type OSUSER, for a use using the crsctl add wallet command, that user is automatical added to this group. This process enables CRS to start user-defined resources as the user that was added to this group.		
		Refer to the <i>Oracle Clusterware Administration and Deployment Guide</i> for details about adding users to a wallet.		
ORA_RAC	SYSRAC privileges for all Oracle Database installations on the server.	The OSRACDBA group for the Windows Operating System. Members of this group have SYSRAC privileges for all Oracle Databases installed on the server.		
ORA_CLIENT_LISTENERS	None	This group is created with service-specific SIDs for Listeners the Client home.		
ORA_HOMENAME_SVCSIDS	None	This group is created with service-specific SIDs for all Service in the DB Client home.		
ORA_GRID_LISTENERS	None	This group is created with Service specific SIDs for all Grid Home Listeners on the system.		



Table 5-1 (Cont.) Operating System Groups Created During Installa

Operating System Group Names	Database Privileges	Description
ORA_INSTALL	None	This group is created with Oracle Home Users for all Oracle homes on the system .
		Virtual accounts for databases and listeners for all virtual account-based homes are added to this group.

During installation, the gridconfig.bat script creates the services and groups on each node of the cluster. The installed files and permissions are owned by the Oracle Installation user, and require the Administrator privilege.

Oracle creates and populates the groups listed in this table during installation to ensure proper operation of Oracle products. You can manually add other users to these groups to assign these database privileges to other Windows users.

Members of the ORA_DBA group can use operating system authentication to administer all Oracle databases installed on the server. Members of the ORA_HOMENAME_DBA group, where HOMENAME is the name of a specific Oracle installation, can use operating system authentication to manage only the databases that run from that Oracle home.

Related Topics

Standard Administration and Job Role Separation User Groups

Example of Using Role-Allocated Groups and Users

You can use role-allocated groups and users that is compliant with an Optimal Flexible Architecture (OFA) deployment.

Assumptions:

- The user installing the Oracle Grid Infrastructure software is named RACDOMAIN\grid. This user was created before starting the installation.
- The name of the home directory for the Oracle Grid Infrastructure installation is OraGrid21c.
- The user installing the Oracle RAC software is named oracle. This user was created before starting the installation.

During installation of Oracle RAC, an Oracle Home user named RACDOMAIN\oradba1 is specified. The oradba1 user is a Windows domain user that was created before the installation was started.

The name of the Oracle home for the Oracle RAC installation is OraRAC21c home1.

You have a second, Oracle Database installation (not Oracle RAC) on this server. The
installation was performed by the oracle user. The Oracle Home user is oradba2, and this user
was not created before starting the installation.

The Oracle Home name is OraDB21c home1.

 Both the Oracle databases and Oracle Clusterware are configured to use Oracle ASM for data storage.

After installing the Oracle software, you have the following groups and users:



Operating System Group Name	Type of Group	Members
ORA_DBA	OSDBA group	oracle, RACDOMAIN\grid, and the Local System built-in Windows account
ORA_OraRAC21c_home1_DBA	OSDBA group for the Oracle RAC home directory	RACDOMAIN\oradba1
ORA_OraDB21c_home1_DBA	OSDBA group for the Oracle Database home directory	oradba2
ORA_OPER	OSOPER group	none
ORA_OraRAC21c_home1_OPER	OSOPER group for the Oracle RAC home directory	none
ORA_OraDB21c_home1_OPER	OSOPER group for the Oracle Database home directory	none
ORA_ASMADMIN	OSASM group	RACDOMAIN\grid and the Local System built-in Windows account, and the database service IDs
ORA_ASMOPER	OSOPER for ASM group	none
ORA_ASMDBA	OSDBA for ASM group for Oracle ASM clients	RACDOMAIN\grid, oracle, the Local System built-in Windows account, and Oracle Home Users of database homes
ORA_RAC21c_home1_SYSBACKUP, ORA_RAC21c_home1_SYSDG, and ORA_RAC21c_home1_SYSKM	Specialized role groups that authorize users with the SYSBACKUP, SYSDG, and SYSKM system privileges.	none
ORA_DB21c_home1_SYSBACKUP, ORA_DB21c_home1_SYSDG, and ORA_DB21c_home1_SYSKM	Specialized role groups that authorize users with the SYSBACKUP, SYSDG, and SYSKM system privileges.	none

If there are no users listed for an operating system group, then that means the group has no members after installation.

Configuring User Accounts

When installing Oracle Grid Infrastructure for a cluster, you run the installer software as an Administrator user.

Before starting the installation, there are a few checks you need to perform for the Oracle Installation users, to ensure the installation will succeed.

- Configuring Environment Variables for the Oracle Installation User
 The installer uses environment variables set for the Oracle Installation User.
- Verifying User Privileges to Update Remote Nodes
 You must ensure that operations that are performed on multiple nodes can be performed during installation of the Oracle Grid Infrastructure software.
- Managing User Accounts with User Account Control
 To ensure that only trusted applications run on your computer, Windows Server provides
 User Account Control.



Configuring Environment Variables for the Oracle Installation User

The installer uses environment variables set for the Oracle Installation User.

 Before starting the Oracle Grid Infrastructure installation, ensure that you have correctly set the %TEMP% environment variable. Also, ensure that you have set the variable to the same value on all cluster nodes.

Related Topics

Checking the Available TEMP Disk Space

Verifying User Privileges to Update Remote Nodes

You must ensure that operations that are performed on multiple nodes can be performed during installation of the Oracle Grid Infrastructure software.

For the installation to be successful, you must use the same user name and password on each node in a cluster or use a domain user. You must explicitly grant membership in the local Administrators group to the installation user on all of the nodes in your cluster.

1. Determine if User Account Control (UAC) remote restrictions have been disabled for the local installation user. If you are using a domain user for installation, then skip this step.

Check the value of the LocalAccountTokenFilterPolicy registry entry for the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System. The value should be set to 1. If this registry entry does not exist, then, do the following:

- a. Click Start, click Run, type regedit, and then press Enter.
- Navigate to the registry key HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Syste m.
- c. On the Edit menu, select New, and then click DWORD Value.
- Type LocalAccountTokenFilterPolicy and then press Enter.
- e. Right-click LocalAccountTokenFilterPolicy, then click **Modify**.
- f. In the Value data box, type 1, then click **OK**.
- g. Exit the registry editor.

By setting the value of LocalAccountTokenFilterPolicy to 1, a user who is a member of the local administrators group on the target remote computer establishes a remote administrative connection with an elevated token, enabling the user to perform administrative tasks. If you do not disable the UAC remote restrictions for administrative users, then when installing Oracle Grid Infrastructure on multiple nodes you might encounter the following error:

INS-40937 The following hostnames are invalid

Before running OUI, from the node where you intend to run the installer, verify that the user account you are using for the installation is configured as a member of the Administrators group on each node in the cluster.

Enter the following command for each node that is a part of the cluster where *nodename* is the node name:

net use \\nodename\C\$



- 3. If you will be using other disk drives in addition to the C: drive, then repeat the net use command for every node in the cluster, substituting the drive letter for each drive you plan to use.
- Verify the installation user is configured to update the Windows registry on each node in the cluster.
 - a. Run regedit from the Run menu or the command prompt.
 - From the File menu select Connect Network Registry.
 - c. In the 'Enter the object name...' edit box enter the name of a remote node in the cluster, then click OK.
 - **d.** Wait for the node to appear in the registry tree.

If the remote node does not appear in the registry tree or you are prompted to fill in a username and password, then you must resolve the permissions issue at the operating system level before proceeding with the Oracle Grid Infrastructure installation.



(i) See Also

Oracle Database Administrator's Reference for Microsoft Windows for information about NTFS file system and Windows registry permissions

Managing User Accounts with User Account Control

To ensure that only trusted applications run on your computer, Windows Server provides User Account Control.

If you have enabled the User Account Control security feature, then depending on how you have it configured, OUI prompts you for either your consent or your credentials when installing Oracle Database. Provide either the consent or your Windows Administrator credentials as appropriate.

You must have Administrator privileges to run some Oracle tools, such as DBCA, NETCA, and OPatch, or to run any tool or application that writes to any directory within the Oracle home. If User Account Control is enabled and you are logged in as the local Administrator, then you can successfully run each of these commands. However, if you are logged in as "a member of the Administrators group," then you must explicitly run these tools with Windows Administrator privileges.

All of the Oracle shortcuts that require Administrator privileges are automatically run as an "Administrator" user when you click the shortcuts. However, if you run the previously mentioned tools from a Windows command prompt, then you must run them from an Administrator command prompt.

OPatch does not have a shortcut and must be run from an Administrator command prompt.

Creating Oracle Software Directories

During installation, you are prompted to provide a path to a home directory to store Oracle Grid Infrastructure software.

You also need to provide a home directory when installing Oracle RAC. Each directory has certain requirements that must be met for the software to work correctly.

Oracle Universal Installer creates the directories during installation if they do not exist.



Temporary Directories

To install properly across all nodes, OUI uses the temporary folders defined within Microsoft Windows.

Oracle Base Directory for the Grid User

During installation, you are prompted to specify an Oracle base location, which is owned by the user performing the installation.

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

Grid Home Directory

When installing Oracle Grid Infrastructure, you must determine the location of the Oracle home for Oracle Grid Infrastructure software (Grid home).

Oracle Inventory Directory

The **Oracle Inventory directory** is the central inventory location for all Oracle software installed on a server.

- Requirements for the Oracle Grid Infrastructure Home Directory
 Review directory path requirements for the Oracle Grid Infrastructure Home directory.
- About Creating the Oracle Base Directory Path
 The Oracle base directory for the Oracle Installation User for Oracle Grid Infrastructure is the location where diagnostic and administrative logs, and other logs associated with Oracle ASM and Oracle Clusterware are stored.

Temporary Directories

To install properly across all nodes, OUI uses the temporary folders defined within Microsoft Windows.

The TEMP and TMP environment variables should point to the same local directory on all nodes in the cluster.

By default, these settings are defined as %USERPROFILE%\Local Settings\Temp and %USERPROFILE%\Local Settings\Tmp in the Environment Settings of My Computer. It is recommended to explicitly redefine these as %WINDIR%\temp and %WINDIR%\temp.

For example, if Windows is installed on the C drive, then the temporary directories would be defined as C:\Windows\temp or C:\Windows\temp for all nodes.

Oracle Base Directory for the Grid User

During installation, you are prompted to specify an Oracle base location, which is owned by the user performing the installation.

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, for an Oracle Grid Infrastructure installation, you must create a different path for the Grid home, so that the path for Oracle base remains available for other Oracle installations.

Multiple Oracle Database installations can use the same Oracle base directory. The Oracle Grid Infrastructure installation uses a different directory path, one outside of Oracle base. If you use different operating system users to perform the Oracle software installations, then each user will have a different default Oracle base location.



After installing Oracle Database 12c release 1 (12.1) (or later) release with a Windows User Account as Oracle Home User, do not install older releases of Oracle Database that share the same Oracle Base Directory. During installation of the software for older releases, the ACLs are reset and Oracle Database 12c release 1 (12.1) (or later) services may not be able to access the Oracle Base directory and files.

In a default Windows installation, the Oracle base directory appears as follows, where X represents a disk drive and username is the name of the installation owner, which is also the name of the software installation owner:

 $X:\langle app \rangle username$

Because you can have only one Oracle Grid Infrastructure installation on a cluster, and because all upgrades are out-of-place upgrades, Oracle recommends that you have an Oracle base for the Oracle Installation user for Oracle Grid Infrastructure (for example, C:\app\grid). You should also use a Grid home for the Oracle Grid Infrastructure binaries using the release number of that installation (for example, C:\app\12.2.0\grid).

During installation, ownership of the path to the Grid home is changed to the LocalSystem, if specified. Using separate paths for the Oracle base directory for the grid user and the Grid home directory means the path for Oracle base remains available for other Oracle software installations by the grid user. If you do not create a unique path to the Grid home, then after the Oracle Grid Infrastructure installation, you might encounter permission errors for other installations, including any existing installations under the same path.

⚠ Caution

For Oracle Grid Infrastructure (for a cluster) installations, note the following restriction for the Oracle Grid Infrastructure home (the Grid home directory for Oracle Grid Infrastructure):

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.

(i) Note

Placing Oracle Grid Infrastructure for a cluster binaries on a cluster file system is not supported.

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 Directory for Grid Infrastructure Management Repository

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

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

C:\app\21.0.0\gimr



(i) Note

The GIMR home directory must be owned by the grid user and you must not create user databases in this directory.

For Grid Infrastructure Management Repository installations, note the following restrictions for the Grid Infrastructure Management Repository binary home (GIMR home directory):

- It must not be placed under one of the Oracle base directories, including the Oracle base directory of the Oracle Grid Infrastructure installation owner.
- It must not be placed in the home directory of an installation owner. These requirements are specific to Grid Infrastructure Management Repository installations.

Grid Home Directory

When installing Oracle Grid Infrastructure, you must determine the location of the Oracle home for Oracle Grid Infrastructure software (Grid home).

The Oracle Home directory that Oracle Grid Infrastructure is installed in is the Grid home. Oracle ASM is also installed in this home directory.

If you plan to install Oracle RAC, you must choose a different directory in which to install the Oracle Database software. The location of the Oracle RAC installation is the Oracle home.

The 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 \path\version\user, where \path is a string constant (C:\).



 \forall version is the version of the software (21.0.0), and \forall user is the installation owner of the Oracle Grid Infrastructure software (grid).

C:\21.0.0\grid



(i) Note

For Oracle Grid Infrastructure for cluster installations, the Grid home must not be placed in the home directory of an installation owner.

Oracle Grid Infrastructure for a standalone server (Oracle Restart) can be installed under the Oracle base for the Oracle Database installation.

Oracle Inventory Directory

The Oracle Inventory directory is the central inventory location for all Oracle software installed on a server.

The first time you install Oracle software on a system, the installer checks to see if an Oracle Inventory directory exists. The location of the Oracle Inventory directory is determined by the Windows Registry key HKEY LOCAL MACHINE\SOFTWARE\Oracle\inst loc. If an Oracle Inventory directory does not exist, then the installer creates one in the default location of C:\Program Files\Oracle\Inventory. The central inventory is created on all cluster nodes.



(i) Note

Changing the value for inst_loc in the Windows registry is not supported.

The Oracle Inventory directory is not installed under the Oracle base directory for the Oracle Installation user. This is because all Oracle software installations share a common Oracle Inventory, so there is only one Oracle Inventory for all users, whereas there is a separate Oracle Base directory for each user.

The Oracle Inventory directory contains the following:

- A registry of the Oracle home directories (Oracle Grid Infrastructure and Oracle Database) on the system
- Installation logs and trace files from installations of Oracle software. These files are also copied to the respective Oracle homes for future reference.
- Other metadata inventory information regarding Oracle installations are stored in the individual Oracle home inventory directories, and are separate from the central inventory.

Requirements for the Oracle Grid Infrastructure Home Directory

Review directory path requirements for the Oracle Grid Infrastructure Home directory.

- It is located in a path outside existing Oracle homes, including Oracle Clusterware homes.
- It is not located in a user home directory.
- You create the path before installation, and then the installer for Oracle Grid Infrastructure creates the directories in the path.



Oracle recommends that you install Oracle Grid Infrastructure on local homes, rather than using a shared home on shared storage.

For installations with Oracle Grid Infrastructure only, Oracle recommends that you create a path compliant with Oracle Optimal Flexible Architecture (OFA) guidelines, so that Oracle Universal Installer (OUI) can select that directory during installation.

(i) Note

Oracle Grid Infrastructure homes can be placed in a local directory on servers, even if your existing Oracle Clusterware home from a prior release is in a shared location.

If you are installing Oracle Grid Infrastructure for a Standalone Server (Oracle Restart), then the home directory for Oracle Restart can be under the Oracle base directory for the Oracle Installation user for Oracle Database.

Related Topics

- Optimal Flexible Architecture File Path Examples
 This topic shows examples of hierarchical file mappings of an Optimal Flexible Architecture-compliant installation.
- Oracle Database Installation Guide for Microsoft Windows

About Creating the Oracle Base Directory Path

The Oracle base directory for the Oracle Installation User for Oracle Grid Infrastructure is the location where diagnostic and administrative logs, and other logs associated with Oracle ASM and Oracle Clusterware are stored.

If the directory or path you specify during installation for the Grid home does not exist, then OUI creates the directory.



Placing Oracle Grid Infrastructure for a cluster binaries on a cluster file system is not supported.

Related Topics

Oracle Base Directory Naming Convention

Configuring Shared Storage for Oracle Database and Grid Infrastructure

Review supported storage options and perform prerequisite tasks as part of your installation planning process.

(i) Note

- If you are currently using OCFS for Windows as your shared storage, then you
 must migrate to using Oracle ASM before the upgrade of Oracle Database and
 Oracle Grid Infrastructure.
- Starting with Oracle Database 21c, the name of Oracle Automatic Storage
 Management Cluster File System (Oracle ACFS) is changed to Oracle Advanced
 Cluster File System (Oracle ACFS).
- Starting with Oracle Database 21c, the Oracle Grid Infrastructure feature Oracle Advanced Cluster File System (Oracle ACFS) is desupported with Microsoft Windows.

⚠ Caution

Oracle ACFS and Oracle Automatic Storage Management Dynamic Volume Manager (Oracle ADVM) will not be accessible after you install or upgrade to Oracle Grid Infrastructure 21c.

- Supported Storage Options for Oracle Grid Infrastructure and Oracle RAC

 Both Oracle Clusterware and the Oracle RAC database use files that must be available to all the nodes in the cluster. The following table shows the storage options supported for storing Oracle Clusterware and Oracle RAC files.
- Storage Considerations for Oracle Grid Infrastructure and Oracle RAC
 For all installations, you must choose the storage option to use for Oracle Grid
 Infrastructure (Oracle Clusterware and Oracle ASM), and Oracle Real Application Clusters
 (Oracle RAC) databases.
- Guidelines for Choosing a Shared Storage Option
 You can choose any combination of the supported shared storage options for each file type if you satisfy all requirements listed for the chosen storage option.
- About Migrating Existing Oracle ASM Instances
 You can use Oracle Automatic Storage Management Configuration Assistant (ASMCA) to upgrade the existing Oracle ASM instance to Oracle ASM 12c release 1 (12.1) or higher.
- Preliminary Shared Disk Preparation
 When using shared storage on a Windows platform, there are additional preinstallation tasks to complete.



Configuring Disk Partitions on Shared Storage
 Use the disk administration tools provided by the operating system or third-party vendors to create disk partitions.

Supported Storage Options for Oracle Grid Infrastructure and Oracle RAC

Both Oracle Clusterware and the Oracle RAC database use files that must be available to all the nodes in the cluster. The following table shows the storage options supported for storing Oracle Clusterware and Oracle RAC files.

Table 6-1 Supported Storage Options for Oracle Clusterware and Oracle RAC Files and Home Directories

Storage Option	OCR and Voting Files	Oracle Grid Infrastructure Home	Oracle RAC Home	Oracle RAC Database Files	Oracle Recovery Files
Oracle Automatic Storage Management (Oracle ASM)	Yes	No	No	Yes	Yes
Direct NFS Client access to a certified network attached storage (NAS) filer	No	No	No	Yes	Yes
Note: NFS or Direct NFS Client cannot be used for Oracle Clusterware files.					
Direct-attached storage (DAS)	No	No	Yes	Yes	Yes
Shared disk partitions (raw devices)	No	No	No	No	No
Local file system (NTFS formatted disk)	No	Yes	Yes	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 only use Oracle ASM to store Oracle Clusterware files.
- Direct use of raw or block devices is not supported. You can only use raw or block devices with Oracle ASM.

Related Topics

Oracle Database Upgrade Guide



Storage Considerations for Oracle Grid Infrastructure and Oracle RAC

For all installations, you must choose the storage option to use for Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM), and Oracle Real Application Clusters (Oracle RAC) databases.

Storage Considerations for Oracle Clusterware

Oracle Clusterware voting files are used to monitor cluster node status, and Oracle Cluster Registry (OCR) files contain configuration information about the cluster. You must store Oracle Cluster Registry (OCR) and voting files in Oracle ASM disk groups. 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.

Storage Considerations for Oracle RAC

Oracle ASM is a supported storage option for database and recovery files. For all installations, Oracle recommends that you create at least two separate Oracle ASM disk groups: One for Oracle Database data files, and one for recovery files. Oracle recommends that you place the Oracle Database disk group and the recovery files disk group in separate failure groups.

- If you do not use Oracle ASM for database files, then Oracle recommends that you place the data files and the Fast Recovery Area in shared storage located outside of the Oracle home, in separate locations, so that a hardware failure does not affect availability.
- You can choose any combination of the supported storage options for each file type provided that you satisfy all requirements listed for the chosen storage options.
- To use Oracle ASM with Oracle RAC, and if you are configuring a new Oracle ASM instance, then your system must meet the following conditions:
 - All nodes on the cluster have Oracle Clusterware and Oracle ASM 21c installed as part of an Oracle Grid Infrastructure for a cluster installation.
 - Any existing Oracle ASM instance on any node in the cluster is shut down.
 - To provide voting file redundancy, one Oracle ASM disk group is sufficient. The Oracle ASM disk group provides three or five copies depending on the redundancy level selected.

You can use NFS, with or without Direct NFS, to store Oracle Database data files. You cannot use NFS as storage for Oracle Clusterware files.

Guidelines for Choosing a Shared Storage Option

You can choose any combination of the supported shared storage options for each file type if you satisfy all requirements listed for the chosen storage option.

- Guidelines for Using Oracle ASM Disk Groups for Storage
 Plan how you want to configure Oracle ASM disk groups for deployment.
- Guidelines for Using Direct Network File System (NFS) with Oracle RAC
 Network-attached storage (NAS) systems use a network file system (NFS) to access data.

 You can store Oracle RAC data files and recovery files on a supported NAS server using Direct NFS Client.



Guidelines for Using Oracle ASM Disk Groups for Storage

Plan how you want to configure Oracle ASM disk groups for deployment.

During Oracle Grid Infrastructure installation, you can create one or two Oracle ASM disk groups. After the Oracle Grid Infrastructure installation, you can create additional disk groups using Oracle Automatic Storage Management Configuration Assistant (ASMCA), SQL*Plus, or Automatic Storage Management Command-Line Utility (ASMCMD).

Choose to create a second disk group during Oracle Grid Infrastructure installation. The first disk group stores the Oracle Cluster Registry (OCR), voting files, and the Oracle ASM password file. The second disk group stores the Grid Infrastructure Management Repository (GIMR) data files and Oracle Cluster Registry (OCR) backup files. Oracle strongly recommends that you store the OCR backup files in a different disk group from the disk group where you store OCR files. In addition, having a second disk group for GIMR is advisable for performance, availability, sizing, and manageability of storage.



Note

You must specify the Grid Infrastructure Management Repository (GIMR) location at the time of installing Oracle Grid Infrastructure. You cannot migrate the GIMR from one disk group to another later.

If you install Oracle Database or Oracle RAC after you install Oracle Grid Infrastructure, then you can either use the same disk group for database files, OCR, and voting files, or you can use different disk groups. If you create multiple disk groups before installing Oracle RAC or before creating a database, then you can do one of the following:

- Place the data files in the same disk group as the Oracle Clusterware files.
- Use the same Oracle ASM disk group for data files and recovery files.
- Use different disk groups for each file type.

If you create only one disk group for storage, then the OCR and voting files, database files, and recovery files are contained in the one disk group. If you create multiple disk groups for storage, then you can place files in different disk groups.

Oracle Database Configuration Assistant (DBCA) does not have the functionality to create disk groups for Oracle ASM. You can use ASMCA to create disk groups.

Related Topics

Oracle Automatic Storage Management Administrator's Guide

Guidelines for Using Direct Network File System (NFS) with Oracle RAC

Network-attached storage (NAS) systems use a network file system (NFS) to access data. You can store Oracle RAC data files and recovery files on a supported NAS server using Direct NFS Client.

NFS file systems must be mounted and available over NFS mounts before you start the Oracle RAC installation. See your vendor documentation for NFS configuration and mounting information.

Note that the performance of Oracle Database software and the databases that use NFS storage depend on the performance of the network connection between the database server



and the NAS device. For this reason, Oracle recommends that you connect the database server (or cluster node) to the NAS device using a private, dedicated, network connection, which should be Gigabit Ethernet or better.

About Migrating Existing Oracle ASM Instances

You can use Oracle Automatic Storage Management Configuration Assistant (ASMCA) to upgrade the existing Oracle ASM instance to Oracle ASM 12c release 1 (12.1) or higher.

ASMCA is located in the path *Grid_home*\bin. You can also use ASMCA to configure failure groups, and Oracle ASM volumes.



(i) Note

You must first shut down all database instances and applications on the node with the existing Oracle ASM instance before upgrading it.

During installation, if you chose to use Oracle ASM and ASMCA detects that there is a prior Oracle ASM release installed in another Oracle ASM home, then after installing the Oracle ASM 21c software, you can start ASMCA to upgrade the existing Oracle ASM instance.

On an existing Oracle Clusterware or Oracle RAC installation, if the prior release of Oracle ASM instances on all nodes is Oracle ASM 12c release 1 or higher, then you are provided with the option to perform a rolling upgrade of Oracle ASM instances. If the prior release of Oracle ASM instances on an Oracle RAC installation are from an Oracle ASM release prior to Oracle ASM 11g release 1, then rolling upgrades cannot be performed. Oracle ASM is then upgraded on all nodes to 21c.

Preliminary Shared Disk Preparation

When using shared storage on a Windows platform, there are additional preinstallation tasks to complete.

- **Disabling Write Caching**
 - You must disable write caching on all disks that will be used to share data between the nodes in your cluster.
- **Enabling Automounting for Windows** Even though the automount feature is enabled by default, you should verify that automount is enabled.

Disabling Write Caching

You must disable write caching on all disks that will be used to share data between the nodes in your cluster.

- Click Start, then select Administrative Tools, then Computer Management, then Device Manager, and then Disk drives.
- Expand the Disk drives and double-click the first drive that will be used by Oracle software.
- Under the Policies tab for the selected drive, uncheck the option that enables write caching.
- 4. Double-click each of the other drives that will be used by Oracle Clusterware and Oracle RAC and disable write caching as described in Step 3.





Any disks that you use to store files, including database files, that will be shared between nodes, must have write caching disabled.

Enabling Automounting for Windows

Even though the automount feature is enabled by default, you should verify that automount is enabled.

You must enable automounting when using:

- Raw partitions for Oracle ASM
- Oracle Clusterware
- Logical drives for Oracle ASM

① Note

Raw partitions are supported only when upgrading an existing installation using the configured partitions. On new installations, using raw partitions is not supported by ASMCA or OUI, but is supported by the software if you perform manual configuration

1. To determine if automatic mounting of new volumes is enabled, use the following commands:

C:\> diskpart

DISKPART> automount

Automatic mounting of new volumes disabled.

- To enable automounting:
 - a. Enter the following commands at a command prompt:

C:\> diskpart

DISKPART> automount enable

Automatic mounting of new volumes enabled.

- **b.** Type exit to end the diskpart session.
- c. Repeat steps 1 and 2 for each node in the cluster.

(i) Note

All nodes in the cluster must have automatic mounting enabled to correctly install Oracle RAC and Oracle Clusterware. Oracle recommends that you enable automatic mounting before creating any logical partitions for use by the database or Oracle ASM.

You must restart each node after enabling disk automounting.

After disk automounting is enabled and the node is restarted, automatic mounting remains active until it is disabled.



Configuring Disk Partitions on Shared Storage

Use the disk administration tools provided by the operating system or third-party vendors to create disk partitions.

You can create the disk partitions using either the Disk Management Interface or the DiskPart utility, both of which are provided by the operating system.

- Creating Disk Partitions Using the Disk Management Interface
 Use the graphical user interface Disk Management snap-in to manage disks.
- Creating Disk Partitions using the DiskPart Utility
 You can also use the DiskPart utility to manage disks.

Creating Disk Partitions Using the Disk Management Interface

Use the graphical user interface Disk Management snap-in to manage disks.

You can create either Primary partitions or Extended partitions using this method.

- To access the Disk Management snap-in, do one of the following:
 - Type diskmgmt.msc at the command prompt to create Primary partitions.
 Use the command diskpart.exe to create Extended partitinos.
 - From the **Start** menu, select **Administrative Tools**, then **Computer Management**. Then select the **Disk Management** node in the Storage tree.
- Create primary partitions and logical drives in extended partitions by selecting the New Simple Volume option.

Use a basic disk with a primary boot record partition style as an extended partition for creating partitions. Do not use spanned volumes or striped volumes. These options convert the volume to a dynamic disk. Oracle Automatic Storage Management does not support dynamic disks.

- a. In the Assign Drive Letter or Path window, select Do not assign a drive letter or drive path
- **b.** In the Format Partition window, select **Do not format this volume** to specify raw partition.
- 3. On each node in the cluster, ensure that the partitions are visible.

In the Disk Management snap-in, from the Action menu, choose **Rescan Disks**.

- 4. If the disks appear with drive letters on the remote node, then remove the drive letter.
 - a. In the Disk Management snap-in, select the disk that has an assigned drive letter.
 - b. Right-click the selected disk and choose Change Drive Letter and Paths...
 - Select the drive letter and click Remove.
 - d. Repeat these steps for each shared disk partition that has an assigned drive letter.

Related Topics

My Oracle Support Note 782450.1

Creating Disk Partitions using the DiskPart Utility

You can also use the DiskPart utility to manage disks.



1. From an existing node in the cluster, run the DiskPart utility as follows:

C:\> diskpart DISKPART>

List the available disks.

By specifying its disk number (n), select the disk on which you want to create a partition.

DISKPART> list disk DISKPART> select disk n

3. Create an extended partition:

DISKPART> create part ext

4. Create a logical drive of the desired size after the extended partition is created using the following syntax:

DISKPART> create part log [size=n] [offset=n] [noerr]

- 5. Repeat steps 2 through 4 for the second and any additional partitions. An optimal configuration is one partition for the Oracle home and one partition for Oracle Database files.
- 6. List the available volumes, and remove any drive letters from the logical drives you plan to use.

DISKPART> list volume DISKPART> select volume n DISKPART> remove

7. On each node in the cluster, ensure that the partitions are visible.

In the Windows Disk Management snap-in, from the Action menu, choose Rescan Disks.

- 8. If the disks appear with drive letters on the remote node, then remove the drive letter.
 - In the Windows Disk Management snap-in, select the disk that has an assigned drive letter.
 - b. Right-click the selected disk and choose Change Drive Letter and Paths...
 - Select the drive letter and click Remove.
 - d. Repeat these steps for each shared disk partition that has an assigned drive letter.

Configuring Storage for Oracle Automatic Storage Management

Identify storage requirements and Oracle Automatic Storage Management (Oracle ASM) disk group options.

When configuring disks for use with Oracle ASM, you can use the asmtool utility to mark the disks prior to installation.

- Identifying Storage Requirements for Using Oracle ASM for Shared Storage
 Before installing Oracle Grid Infrastructure, you must identify and determine how many
 devices are available for use by Oracle ASM, the amount of free disk space available on
 each disk, and the redundancy level to use with Oracle ASM.
- 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, and installing the starter database, for various Oracle Cluster deployments.
- About the Grid Infrastructure Management Repository
 The Grid Infrastructure Management Repository (GIMR) is a multitenant database with a pluggable database (PDB) for the GIMR of each cluster.
- Restrictions for Disk Partitions Used By Oracle ASM
 Be aware of the following restrictions when configuring disk partitions for use with Oracle ASM:
- Preparing Your System to Use Oracle ASM for Shared Storage
 To use Oracle ASM as the shared storage solution for Oracle Clusterware or Oracle RAC files, you must perform certain tasks before you begin the software installation.
- Marking Disk Partitions for Oracle ASM Before Installation
 The only partitions that OUI displays for Windows systems are logical drives that are on disks and have been marked (or stamped) with asmtoolg or by Oracle Automatic Storage Management (Oracle ASM) Filter Driver.

Identifying Storage Requirements for Using Oracle ASM for Shared Storage

Before installing Oracle Grid Infrastructure, you must identify and determine how many devices are available for use by Oracle ASM, the amount of free disk space available on each disk, and the redundancy level to use with Oracle ASM.

When Oracle ASM provides redundancy, you must have sufficient capacity in each disk group to manage a re-creation of data that is lost after a failure of one or two failure groups.





As you progress through the following steps, make a list of the raw device names you intend to use to create the Oracle ASM disk groups and have this information available during the Oracle Grid Infrastructure installation or when creating your Oracle RAC database.

- Plan your Oracle ASM disk groups requirement. If you choose to store the Grid Infrastructure Management Repository (GIMR) on a separate Oracle ASM disk group, then you require two separate Oracle ASM disk groups, one for OCR and voting files and the other for the GIMR.
- Determine whether you want to use Oracle ASM for Oracle Database files, recovery files, or all file types. Oracle Database files include data files, control files, redo log files, the server parameter file, and the password file.

Note

- You do not have to use the same storage mechanism for Oracle Database files and recovery files. You can use a shared file system for one file type and Oracle ASM for the other.
- There are two types of Oracle Clusterware files: OCR files and voting files. You must use Oracle ASM to store OCR and voting files.
- 3. Choose the Oracle ASM redundancy level to use for the Oracle ASM disk group.

The redundancy level that you choose for the Oracle ASM disk group determines how Oracle ASM mirrors files in the disk group, and determines the number of disks and amount of disk space that you require. 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 stores the Oracle Clusterware voting files. The quorum failure group ensures 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 determines 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 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.

(i) 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.

Determine the total amount of disk space that you require for the Oracle Clusterware files using Oracle ASM for shared storage.

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.

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 redundancy disk groups, the default AU size is 1 MB.

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 nodes = Number of nodes in cluster clients - Number of database instances for each node disks - Number of disks in disk group

For example, for a four-node Oracle RAC installation, using three disks in a normal redundancy disk group, you require an additional 5293 MB of space: [2 * 4 * 3] + [2 * (4 * (4*(4+1)+30)+(64*4)+533)] = 5293 MB

7. Optionally, identify failure groups for the Oracle ASM disk group devices.



(i) Note

You only have to complete this step if you plan to use an installation method that includes configuring Oracle ASM disk groups before installing Oracle RAC, or creating an Oracle RAC database.

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. Failure groups define Oracle ASM disks that share a common potential failure mechanism. By default, each device comprises its own failure group.



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 controller fails. The HBA in this example is a single point of failure. To protect against failures of this type, you could use two HBA fabric paths, each with two disks, and define a failure group for the disks attached to each HBA. This configuration enables the disk group to tolerate the failure of one HBA fabric path.

① Note

You can define custom failure groups during installation of Oracle Grid Infrastructure. You can also define failure groups after installation of Oracle Grid Infrastructure 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 the grid 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.

Related Topics

About Oracle Extended Clusters

An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

(i) See Also

- Oracle Automatic Storage Management Administrator's Guide for more information about Oracle ASM file, quota, and failure groups
- Storage Checklist for Oracle Grid Infrastructure
- 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.



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, and installing the starter database, for various Oracle Cluster deployments.

Total Storage Space for Database Files Required by Redundancy Type

The following table lists the space requirements for Oracle RAC Database data files.

Table 7-1 Oracle ASM Disk Space Requirements for a Multitenant Container Database (CDB) with One Pluggable Database (PDB)

Redundancy Level	Minimum number of disks	Data Files	Recovery Files	Both File Types
External	1	4.5 GB	12.9 GB	17.4 GB
Normal	2	8.6 GB	25.8 GB	34.4 GB
High	3	12.9 GB	38.7 GB	51.6 GB
Flex	3	12.9 GB	38.7 GB	51.6 GB

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

If you create a disk group as part of the installation to install the OCR and voting files, then the installer requires that you create these files on a disk group with at least 2 GB of available space.

Based on the cluster configuration you want to install, the Oracle Clusterware space requirements vary for different redundancy levels. The following table lists the available storage space requirements for external redundancy level for each cluster configuration:



(i) 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.

Table 7-2 Minimum Available Space Requirements for Oracle Standalone Cluster With **GIMR Configuration**

Redunda ncy 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: 10 GB	2 GB	60 GB



Table 7-2 (Cont.) Minimum Available Space Requirements for Oracle Standalone Cluster With GIMR Configuration

Redunda ncy Level	DATA Disk Group	MGMT Disk Group	Oracle Fleet Patching and Provisioning	Total Storage
High/Flex/	3 GB	84 GB	3 GB	90 GB
Extended		Each node beyond four: 15 GB		

Related Topics

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

About the Grid Infrastructure Management Repository

The Grid Infrastructure Management Repository (GIMR) is a multitenant database with a pluggable database (PDB) for the GIMR of each cluster.

Every Oracle Domain Services Cluster contains a Grid Infrastructure Management Repository (GIMR). However, GIMR configuration is optional for Oracle Standalone 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 (QoSM) collects
- Guidelines for the Grid Infrastructure Management Repository
 Review these guidelines while configuring the Grid Infrastructure Management Repository
 (GIMR) for your deployment.

Guidelines for the Grid Infrastructure Management Repository

Review these guidelines while configuring the Grid Infrastructure Management Repository (GIMR) for your deployment.

- Starting with Oracle Grid Infrastructure 19c, configuring the 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).
- Configuring the GIMR is optional for only new installations and does not apply to existing installations.
- If you have a GIMR in your current deployment, then the GIMR is upgraded with Oracle Grid Infrastructure 21c. However, if you do not have a GIMR in your current deployment, then you cannot add the GIMR as part of upgrade to Oracle Grid Infrastructure 21c.

For deployments where there is no GIMR, you can add a GIMR after installing Oracle Grid Infrastructure 21c or after upgrading to Oracle Grid Infrastructure 21c.



Restrictions for Disk Partitions Used By Oracle ASM

Be aware of the following restrictions when configuring disk partitions for use with Oracle ASM:

- With x64 Windows, you can create up to 128 primary partitions for each disk.
- Oracle recommends that you limit the number of partitions you create on a single disk to prevent disk contention. Therefore, you may prefer to use extended partitions rather than primary partitions.

Preparing Your System to Use Oracle ASM for Shared Storage

To use Oracle ASM as the shared storage solution for Oracle Clusterware or Oracle RAC files, you must perform certain tasks before you begin the software installation.

- 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.
- <u>Selecting Disks to use with Oracle ASM Disk Groups</u>
 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.

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.



NOLE

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



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 diskpart.exe.
- 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.

Related Topics

- Preliminary Shared Disk Preparation
- Configuring Disk Partitions on Shared Storage

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, ASMLib libraries, and whether NVMe over Fabrics disks or 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. In addition, if your installation uses multipathing software, then the software might place pseudo-devices in a path that is different from the operating system default.

(i) 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.



Marking Disk Partitions for Oracle ASM Before Installation

The only partitions that OUI displays for Windows systems are logical drives that are on disks and have been marked (or **stamped**) with asmtoolg or by Oracle Automatic Storage Management (Oracle ASM) Filter Driver.

If you chose not to use Oracle ASM Filter Driver (Oracle ASMFD) for configuring and marking the disks to use with Oracle ASM, then you must create disk partitions and use the asmtoolg utility to mark the disk partitions prior to installing Oracle Grid Infrastructure. You also have the option of using the asmtoolg utility during Oracle Grid Infrastructure for a cluster installation.

All disk names created by asmtoolg or asmtool begin with the prefix ORCLDISK followed by a user-defined prefix (the default is DATA), and by a disk number for identification purposes. You can use them as raw devices in the Oracle ASM instance by specifying a name \\.\ORCLDISKprefixn, where prefixn either can be DATA, or a value you supply, and where n represents the disk number.

The asmtoolg and asmtool utilities only work on partitioned disks; you cannot use Oracle ASM on unpartitioned disks. You can also use these tools to reconfigure the disks after installation. These utilities are installed automatically as part of Oracle Grid Infrastructure.



If user account control (UAC) is enabled, then running asmtoolg or asmtool requires administrator-level permissions.

- Using asmtoolg to Mark Disks
 - Use asmtoolg (GUI version) to create device names; use asmtoolg to add, change, delete, and examine the devices available for use in Oracle ASM.
- <u>Using asmtoolg to Remove Disk Stamps</u>
 You can use asmtoolg (GUI version) to delete disk stamps.
- asmtool Command Line Reference
 asmtool is a command-line interface for marking (or stamping) disks to be used with Oracle
 ASM.

Using asmtoolg to Mark Disks

Use asmtoolg (GUI version) to create device names; use asmtoolg to add, change, delete, and examine the devices available for use in Oracle ASM.

- 1. In directory where you unzipped the Oracle Grid Infrastructure image files, go the bin\asmtool folder and double-click asmtoolg.
 - If user access control (UAC) is enabled, then you must create a desktop shortcut to a command window. Right-click the shortcut, select **Run as Administrator**, and launch asmtoolg.
- Select the Add or change label option, and then click Next.
 - asmtoolg shows the devices available on the system. Unrecognized disks have a status of "Candidate device", stamped disks have a status of "Stamped ASM device," and disks that have had their stamp deleted have a status of "Unstamped ASM device." The tool also



shows disks that are recognized by Windows as a file system (such as NTFS). These disks are not available for use as Oracle ASM disks, and cannot be selected.



(i) Note

Dynamic disks are not supported for use with Oracle ASM.

On the Stamp Disks window, select the disks to you want to use with Oracle ASM.

For ease of use, Oracle ASM can generate unique stamps for all of the devices selected for a given prefix. The stamps are generated by concatenating a number with the prefix specified. For example, if the prefix is DATA, then the first Oracle ASM link name is ORCLDISKDATA0.

You can also specify the stamps of individual devices.

- Optional: Select a disk to edit the individual stamp (Oracle ASM link name).
- Click Next.
- Click Finish.

Related Topics

Configuring Disk Partitions on Shared Storage Use the disk administration tools provided by the operating system or third-party vendors to create disk partitions.

Using asmtoolg to Remove Disk Stamps

You can use asmtoolg (GUI version) to delete disk stamps.

- In directory where you unzipped the Oracle Grid Infrastructure image files, go the *Grid_home*\bin folder and double-click asmtoolg.
 - If user access control (UAC) is enabled, then you must create a desktop shortcut to a command window. Open the command window using the Run as Administrator, rightclick the context menu, and launch asmtoolg.
- Select the **Delete labels** option, then click **Next**.
 - The delete option is only available if disks exist with stamps. The delete screen shows all stamped Oracle ASM disks.
- On the Delete Stamps screen, select the disks to unstamp.
- Click Next.
- Click Finish.

asmtool Command Line Reference

asmtool is a command-line interface for marking (or stamping) disks to be used with Oracle ASM.



Option	Description	Example
-add	Adds or changes stamps. You must specify the hard disk, partition, and new stamp name. If the disk is a raw device or has an existing Oracle ASM stamp, then you must specify the -force option.	asmtool -add [-force] \Device\Harddisk1\Partition1 ORCLDISKASM0 \Device\Harddisk2\Partition1 ORCLDISKASM2
-addprefix	Adds or changes stamps using a common prefix to generate stamps automatically. The stamps are generated by concatenating a number with the prefix specified. If the disk is a raw device or has an existing Oracle ASM stamp, then you must specify the force option.	asmtool -addprefix ORCLDISKASM [-force] \Device\Harddisk1\Partition1 \Device\Harddisk2\Partition1
-create	Creates an Oracle ASM disk device from a file instead of a partition. Note: Usage of this command is not supported for production environments.	asmtool -create \\server\share\file 1000 asmtool -create D:\asm\asmfile02.asm 240
-list	List available disks. The stamp, windows device name, and disk size in MB are shown.	asmtool -list
-delete	Removes existing stamps from disks.	asmtool -delete ORCLDISKASM0 ORCLDISKASM1

If user access control (UAC) is enabled, then you must create a desktop shortcut to a command window. Open the command window using the Run as Administrator, right-click the context menu, and launch asmtool.



(i) Note

If you use -add, -addprefix, or -delete, asmtool notifies the Oracle ASM instance on the local node and on other nodes in the cluster, if available, to rescan the available disks.

Related Topics

Configuring Disk Partitions on Shared Storage

Configuring Direct NFS Client for Oracle RAC Data Files

Direct NFS Client is an interface for NFS systems provided by Oracle.

About Direct NFS Client Storage

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.

Creating an oranfstab File for Direct NFS Client

If you use Direct NFS Client, then you must create a configuration file, oranfstab, to specify the options, attributes, and parameters that enable Oracle Database to use Direct NFS Client.

Configurable Attributes for the oranfstab File

You can configure various settings in the oranfstab file.

Mounting NFS Storage Devices with Direct NFS Client

Direct NFS Client determines mount point settings for NFS storage devices based on the configuration information in oranfstab. Direct NFS Client uses the first matching entry as the mount point.

Specifying Network Paths for a NFS Server

Direct NFS Client can use up to four network paths defined in the oranfstab file for an NFS server.

Enabling Direct NFS Client

You enable Direct NFS Client by running the enable_dnfs.bat command.

Performing Basic File Operations Using the ORADNFS Utility

ORADNFS is a utility which enables the database administrators to perform basic file operations over Direct NFS Client on Microsoft Windows platforms.

Monitoring Direct NFS Client Usage

You use data dictionary view to monitor the Direct NFS client.

Disabling Direct NFS Client

You disable Direct NFS Client by running the disable dnfs.bat command.

About Direct NFS Client Storage

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 and NFSv4.1 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 (12.2), Windows Direct NFS Client supports all widely accepted NFS path formats including UNIX-style NFS paths and NFS Version 4 protocol.

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 tunes itself to make optimal use of available resources and enables the storage of data files on supported NFS servers. Direct NFS Client obtains NFS mount points from the oranfstab file.



(i) Note

Use NFS servers supported for Oracle RAC. Check My Oracle Support, as described in Checking Hardware and Software Certification on My Oracle Support.

Direct NFS Client Requirements

- Direct NFS cannot provide service to NFS servers with write size values (wtmax) less than 32768.
- The Oracle files resident on the NFS server that are accessed by Direct NFS Client can also be accessed through a third-party NFS client. The volume must be mounted through Common Internet File System (CIFS) or kernel NFS to enable regular windows utilities and commands, such as copy, and so on, access the database files in the remote location.
- Volumes mounted through CIFS can not be used for storing Oracle database files without configuring Direct NFS Client. The atomic write requirements needed for database writes are not guaranteed through the CIFS protocol, consequently CIFS can only be used for OS level access, for example, for commands such as copy.
- To enable Oracle Database to use Direct NFS Client, the NFS file systems must be mounted and available before you start installation. Direct NFS Client manages settings after installation.
 - If Oracle Database cannot open an NFS server using Direct NFS Client, then an informational message is logged into the Oracle alert log. A trace file is also created, indicating that Direct NFS Client could not connect to an NFS server.
- Some NFS file servers require NFS clients to connect using reserved ports. If your filer is running with reserved port checking, then you must disable it for Direct NFS to operate. To disable reserved port checking, consult your NFS file server documentation.
- For NFS servers that restrict port range, you can use the insecure option to enable clients other than an Administrator user to connect to the NFS server. Alternatively, you can disable Direct NFS Client.
- You can have only one active Direct NFS Client implementation for each instance. Using Direct NFS Client on an instance will prevent another Direct NFS Client implementation.



(i) See Also

- Oracle Database Reference for information on setting the enable_dnfs_dispatcher parameter in the initialization parameter file to enable Direct NFS dispatcher
- Oracle Database Administrator's Guide for guidelines for managing Oracle data files created with Direct NFS client.
- 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 regarding managing Oracle Database data files created with Direct NFS Client or kernel NFS

Creating an oranfstab File for Direct NFS Client

If you use Direct NFS Client, then you must create a configuration file, oranfstab, to specify the options, attributes, and parameters that enable Oracle Database to use Direct NFS Client.

Direct NFS Client looks for the mount point entries in oranfstab. It uses the first matched entry as the mount point.

 Create an oranfstab file and specify configurable attributes for each NFS server that the Direct NFS Client accesses.

For a 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.

The mount point specified in the oranfstab file represents the local path where the database files would reside normally, as if Direct NFS Client was not used. For example, if the location for the data files is C:\app\oracle\oradata\orcl and Direct NFS Client is not used, then specify C:\app\oracle\oradata\orcl as NFS virtual mount point in the corresponding oranfstab file.

2. Save the file to the *Oracle_home*\dbs directory.

When the oranfstab file is placed in *Oracle_home*\dbs, the entries in the file are specific to a single database.

3. If you have a nonshared Oracle home, copy the oranfstab file to the *Oracle_home*\dbs directory on all nodes in the cluster.

All instances that use the shared Oracle home use the same *Oracle_home*\dbs\oranfstab file.

For a nonshared Oracle home, you must keep the oranfstab file synchronized on all the nodes.



① Note

- Direct NFS Client ignores a uid or gid value of 0. The exported path from the NFS server must be accessible for read/write/execute by the user with the uid, gid specified in oranfstab. If neither uid nor gid is listed, then the exported path must be accessible by the user with uid:65534 and gid:65534.
- If you remove an NFS path from oranfstab, you must make the change in all oranfstab
 files used by the Oracle RAC database. Then, you must restart the database for
 the change to be effective. The mount point that you use for the file system must
 be identical on each node.

The following examples show three possible NFS server entries in oranfstab. A single oranfstab can have multiple NFS server entries.

Example 8-1 oranfstab File Using Local and Path NFS Server Entries

The following example of an oranfstab file shows an NFS server entry, where the NFS server, MyDataServer1, uses two network paths specified with IP addresses.

server: MyDataServer1 local: 192.0.2.0 path: 192.0.2.1 local: 192.0.100.0 path: 192.0.100.1 nfs_version: nfsv3

export: /vol/oradata1 mount: C:\APP\ORACLE\ORADATA\ORADATA1

Example 8-2 oranfstab File Using Network Names in Place of IP Addresses, with Multiple Exports, management and community

The following example of an oranfstab file shows an NFS server entry, where the NFS server, MyDataServer2, uses four network paths specified by the network interface to use, or the network connection name. Multiple export paths are also used in this example.

server: MyDataServer2 local: LocalInterface1 path: NfsPath1 local: LocalInterface2 path: NfsPath2 local: LocalInterface3 path: NfsPath3 local: LocalInterface4 path: NfsPath4 nfs_version: nfsv4

export: /vol/oradata2 mount: C:\APP\ORACLE\ORADATA\ORADATA2 export: /vol/oradata3 mount: C:\APP\ORACLE\ORADATA\ORADATA3

management: MgmtPath1 community: private



Example 8-3 oranfstab File Using Kerberos Authentication with Direct NFS Export

In this example, when specified, the security parameter overrides the value of the security_default parameter.

server: nfsserver local: 198.51.100.02 path: 10.0.0.0 local: 198.51.100.03 path: 10.0.0.3

 $export: /private/oracle1/logs \ mount: C:\APP\ORACLE\ORADATA\logs \ security: krb5 \\ export: /private/oracle1/data \ mount: C:\APP\ORACLE\ORADATA\data \ security: krb5p \\ export: /private/oracle1/archive mount: C:\APP\ORACLE\ORADATA\archive \ security: sys$

export: /private/oracle1/data1 mount: C:\APP\ORACLE\ORADATA\data1

security_default: krb5i

Related Topics

Enabling Direct NFS Client
You enable Direct NFS Client by running the enable_dnfs.bat command.

Configurable Attributes for the oranfstab File
 You can configure various settings in the oranfstab file.

Configurable Attributes for the oranfstab File

You can configure various settings in the oranfstab file.

Table 8-1 Configurable Attributes for the oranfstab File

Attribute	Description	
server	The NFS server name.	
path	Up to four network paths to the NFS server, specified either by internet protocol (IP) address, or by name, as displayed using the ifconfig command on the NFS server.	
local	Up to 4 paths on the database host, specified by IP address or by name, as displayed using the ipconfig command on the database host.	
export	The exported path from the NFS server. Use a UNIX-style path.	
mount	The corresponding local mount point for the exported volume. The path can use a Windows-style path, or the following path formats:	
	 nfs://server/export/path/file 	
	 server:/export/path/file 	
	• //server/export/path/file	
mnt_timeout	(Optional) The time (in seconds) for which Direct NFS Client should wait for a successful mount before timing out. The default timeout is 10 minutes (600).	
uid	(Optional) The UNIX user ID to be used by Direct NFS Client to access all NFS servers listed in oranfstab. The default value is uid:65534, which corresponds to user:nobody on the NFS server.	
gid	(Optional) The UNIX group ID to be used by Direct NFS Client to access all NFS servers listed in oranfstab. The default value is gid:65534, which corresponds to group:nogroup on the NFS server.	



Table 8-1 (Cont.) Configurable Attributes for the oranfstab File

Attribute	Description	
nfs_version	(Optional) The NFS protocol version used by the Direct NFS Client. Possible values are NFSv3, NFSv4, and NFSv4.1, and pNFS. The default version is NFSv3. To specify NFSv4.x, then you must configure the nfs_version parameter accordingly in the oranfstab file. Specify nfs_version as pNFS, if you want to use Direct NFS with Parallel NFS.	
security_default	(Optional) The default security mode applicable for all the exported NFS server paths for a server entry. sys is the default value. See the description of the security parameter for the supported security levels for the security_default parameter.	
security	(Optional) The security level, if you want to enable security using Kerberos authentication protocol with Direct NFS Client. This 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.	
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 server.	
community	The community string for use in SNMP queries. Default value is public.	

Related Topics

• Creating an oranfstab File for Direct NFS Client

If you use Direct NFS Client, then you must create a configuration file, oranfstab, to specify the options, attributes, and parameters that enable Oracle Database to use Direct NFS Client.



"Limiting Asynchronous I/O in NFS Server Environments" in *Oracle Database Performance Tuning Guide*



Mounting NFS Storage Devices with Direct NFS Client

Direct NFS Client determines mount point settings for NFS storage devices based on the configuration information in oranfstab. Direct NFS Client uses the first matching entry as the mount point.

If Oracle Database cannot open an NFS server using Direct NFS Client, then an error message is written into the Oracle alert and trace files indicating that Direct NFS Client could not be established.

(i) Note

You can have only one active Direct NFS Client implementation for each instance. Using Direct NFS Client on an instance will prevent another Direct NFS Client implementation.

Direct NFS Client requires an NFS server supporting NFS read/write buffers of at least 16384 bytes.

Direct NFS Client issues writes at wtmax granularity to the NFS server. Direct NFS Client does not serve an NFS server with a wtmax less than 16384. Oracle recommends that you use the value 32768.



(i) See Also

"Supported Storage Options for Oracle Grid Infrastructure and Oracle RAC" for a list of the file types that are supported with Direct NFS Client.

Specifying Network Paths for a NFS Server

Direct NFS Client can use up to four network paths defined in the oranfstab file for an NFS server.

Direct NFS Client performs load balancing across all specified paths. If a specified path fails, then Direct NFS Client reissues all outstanding requests over any remaining paths.



(i) Note

You can have only one active Direct NFS Client implementation for each instance. Using Direct NFS Client on an instance prevents the use of another Direct NFS Client implementation.



(i) See Also

<u>Creating an oranfstab File for Direct NFS Client</u> for examples of configuring network paths for Direct NFS Client attributes in an oranfstab file.



Enabling Direct NFS Client

You enable Direct NFS Client by running the enable_dnfs.bat command.

- Shut down the Oracle Database instance.
- 2. Create and configure an oranfstab file in *ORACLE_BASE_HOME*\dbs.
- 3. Set ORACLE_HOME to Oracle home for which the Direct NFS Client must be enabled.
- **4.** Change directory to *ORACLE_HOME*\bin.
- Run the batch file enable_dnfs.bat.
- 6. Restart the Oracle Database instance.

For Oracle RAC, repeat the above procedure on all nodes in the cluster.

Related Topics

Creating an oranfstab File for Direct NFS Client

If you use Direct NFS Client, then you must create a configuration file, oranfstab, to specify the options, attributes, and parameters that enable Oracle Database to use Direct NFS Client.

Performing Basic File Operations Using the ORADNFS Utility

ORADNFS is a utility which enables the database administrators to perform basic file operations over Direct NFS Client on Microsoft Windows platforms.

ORADNFS is a multi-call binary, which is a single binary that acts like many utilities.

You must be a member of the local ORA_DBA group to use ORADNFS. A valid copy of the oranfstab configuration file must be present in *Oracle_home*\dbs for ORADNFS to operate.

 To execute commands using ORADNFS you issue the command as an argument on the command line.

The following command prints a list of commands available with ORADNES:

C:\> oradnfs help

To display the list of files in the NFS directory mounted as C:\ORACLE\ORADATA, use the following command:

C:\> oradnfs ls C:\ORACLE\ORADATA\ORCL

Monitoring Direct NFS Client Usage

You use data dictionary view to monitor the Direct NFS client.

- Use the following global dynamic performance views for managing Direct NFS Client usage with your Oracle RAC database:
 - GV\$DNFS_SERVERS: Lists the servers that are accessed using Direct NFS Client.
 - GV\$DNFS FILES: Lists the files that are currently open using Direct NFS Client.



- GV\$DNFS_CHANNELS: Lists the open network paths, or channels, to servers for which Direct NFS Client is providing files.
- GV\$DNFS_STATS: Lists performance statistics for Direct NFS Client.

Disabling Direct NFS Client

You disable Direct NFS Client by running the disable_dnfs.bat command.

- 1. Shut down the Oracle Database instance.
- 2. Remove the oranfstab configuration file from ORACLE_BASE_HOME\dbs.
- 3. Set ORACLE_HOME to Oracle home for which the Direct NFS Client must be disabled.
- 4. Change directory to *ORACLE_HOME*\bin.
- 5. Run the batch file, disable dnfs.bat.
- 6. Restart the Oracle Database instance.

For Oracle RAC, repeat the above procedure on all nodes in the cluster.

Installing Oracle Grid Infrastructure for a Cluster

Review this information for installation and deployment options for Oracle Grid Infrastructure.

Oracle Grid Infrastructure consists of Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM). If you plan afterward to install Oracle Database with Oracle Real Application Clusters (Oracle RAC), then this is phase one of a two-phase installation.

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. You may use the Graphical User Interface (GUI) mode, or the silent mode to set up Oracle Grid Infrastructure.

- About Image-Based Oracle Grid Infrastructure Installation
 Installation and configuration of Oracle Grid Infrastructure software is simplified with image-based installation.
- <u>Setup Wizard Installation Options for Creating Images</u>
 Gold image-creation options to use with your Oracle Database or Oracle Grid Infrastructure installation setup wizards.
- Understanding Cluster Configuration Options
 Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 21c.
- About Default File Permissions Set by Oracle Universal Installer
 Oracle Grid Infrastructure is installed in the ORACLE_HOME directory, by default. Oracle
 Universal Installer sets the permissions for this directory, and for all files and directories
 under this directory.
- Installing Oracle Grid Infrastructure for a New Cluster
 Complete this procedure to install Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM) on your cluster.
- 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 Software Binaries for Oracle Grid Infrastructure for a Cluster
 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.
- Configuring the 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 using the Grid Setup Wizard utility (setup.exe).



Setting Ping Targets for Network Checks

Receive notification about the status of public networks by setting the Ping_Targets parameter during the Oracle Grid Infrastructure installation.

Confirming Oracle Clusterware Function

After installation, use the crsctl utility to verify 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
 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.

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).

① Note

You must extract the image software into the directory where you want your Grid home to be located, and then run the <code>%ORACLE_HOME%\setup.exe</code> script to start the Oracle Grid Infrastructure Setup Wizard. Ensure that the Grid home directory path you create is in compliance with the Oracle Optimal Flexible Architecture recommendations.



Setup Wizard Installation Options for Creating Images

Gold image-creation options to use with your Oracle Database or Oracle Grid Infrastructure installation setup wizards.

Options

In image-based installations, you can start your Oracle Database installation or Oracle Grid Infrastructure installation by running the setup wizard setup.exe. This wizard comes with the following image-creation options:



Note

setup.exe is the recommended setup wizard for installing both Oracle Database and Oracle Grid Infrastructure.

Table 9-1 Image-Creation Options for Setup Wizard

Option	Description
-createGoldImage	Creates a gold image from the current Oracle home.
-destinationLocation	Specify the complete path, or location, where the gold image will be created.
-exclFiles	Specify the complete paths to the files to be excluded from the newly created gold image.
—help	Displays help for all the available options.

For example:

cd ORACLE_HOME

setup.exe -createGoldImage -destinationLocation c:\my_db_images -exclFiles $C:\app\oracle\product\21.0.0\dbhome_1\relnotes$

cd Grid home

setup.exe -createGoldImage -destinationLocation c:\my grid images -exclFiles *C*:\app\oracle\product\21.0.0\dbhome_1\relnotes

Where:

c:\my_db_images is the file location where the image zip file is created.

c:\my_grid_images is the file location where the image zip file is created.

C:\app\oracle\product\21.0.0\dbhome_1\relnotes is the file to be excluded from the newly created gold image.

Understanding Cluster Configuration Options

Review these topics to understand the cluster configuration options available in Oracle Grid Infrastructure 21c.



About Oracle Standalone Clusters

An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

About Oracle Extended Clusters

An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

About Oracle Standalone Clusters

An Oracle Standalone Cluster hosts all Oracle Grid Infrastructure services and Oracle ASM locally and requires direct access to shared storage.

The number of cluster nodes in an Oracle Standalone Cluster can be as many as 64 and 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 cluster nodes, with an Oracle ASM instance available to all cluster nodes.

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).

When you deploy an Oracle Standalone Cluster, you can also choose to configure it as an Oracle Extended cluster. An Oracle Extended Cluster consists of nodes that are located in multiple locations or sites.

About Oracle Extended Clusters

An Oracle Extended Cluster consists of nodes that are located in multiple locations called sites.

When you deploy an Oracle Standalone Cluster, you can also choose to configure the cluster as an Oracle Extended Cluster. You can extend an Oracle RAC cluster across two, or more, geographically separate sites, each equipped with its own storage. In the event that one of the sites fails, the other site acts as an active standby.

Both Oracle ASM and the Oracle Database stack, in general, are designed to use enterpriseclass 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 9-2 Oracle ASM Disk Group Redundancy Levels for Oracle Extended Clusters with 2 Data Sites

Redundancy Level	Number of Failure Groups for OCR and Voting Files Disk Groups	Number of Failure Groups for OCR Backup and GIMR Disk Groups
Normal redundancy	1 failure group per data site, 1 quorum failure group	1 failure group per data site
Flex redundancy	1 failure group per data site, 1 quorum failure group	1 failure group per data site, 1 quorum failure group
Extended redundancy	3 failure groups each for 2 data sites, 1 quorum failure group outside the 2 data sites	3 failure groups each for 2 data sites, 1 quorum failure group outside the 2 data sites
High redundancy	Not supported	Not supported

Related Topics

- Identifying Storage Requirements for Using Oracle ASM for Shared Storage
 Before installing Oracle Grid Infrastructure, you must identify and determine how many devices are available for use by Oracle ASM, the amount of free disk space available on each disk, and the redundancy level to use with Oracle ASM.
- Oracle Clusterware Administration and Deployment Guide

About Default File Permissions Set by Oracle Universal Installer

Oracle Grid Infrastructure is installed in the ORACLE_HOME directory, by default. Oracle Universal Installer sets the permissions for this directory, and for all files and directories under this directory.

For the ORACLE_HOME of Oracle Grid Infrastructure, the following permissions are granted:

- Full control Administrators, SYSTEM, ORA_GRID_LISTENERS, Oracle Installation User, Oracle Home User
- Read, Execute, and List Contents Authenticated Users

Related Topics

Oracle Database Administrator's Reference for Microsoft Windows

Installing Oracle Grid Infrastructure for a New Cluster

Complete this procedure to install Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM) on your cluster.

If you have any questions during installation, then 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 installation.

- Log in to Windows as the installation user for Oracle Grid Infrastructure, which must be a member of the Administrators users group.
- Create the Grid home directory.



For example:

md C:\app\21.0.0\grid

Ensure that the Grid home directory path you create is in compliance with the Oracle Optimal Flexible Architecture recommendations.

3. Download the Oracle Grid Infrastructure image files and copy the files to the Grid home.

For example:

cd C:\app\21.0.0\grid unzip -q download_location\grid_home.zip

Unzip the installation image files only in this Grid home directory that you created.



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.

4. Start the Oracle Grid Infrastructure Setup Wizard by running the following command:

Grid_home\setup.exe

You can run this command from a Virtual Network Computing (VNC) session, or Terminal Services in console mode.

- Select one of the following configuration options:
 - Configure Oracle Grid Infrastructure for a New Cluster

Select this option to configure a new Oracle Flex Cluster deployment.

Configure Oracle Grid Infrastructure for a standalone server (Oracle Restart)

Select this option to install Oracle Grid Infrastructure in an Oracle Restart configuration. Use this option for single servers supporting Oracle Database and other applications.

Upgrade Oracle Grid Infrastructure

Select this option to upgrade Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).

Set Up Software Only

Select this option to configure Oracle Grid Infrastructure software and register the Oracle Grid Infrastructure home with the central inventory

- **6.** Choose the type of Oracle Flex Cluster from the following options:
 - Configure an Oracle Standalone Cluster: Select this option to install Oracle Grid Infrastructure software for Oracle Standalone Cluster.
 - Configure an Oracle Domain Services Cluster: Select this option to install Oracle Grid Infrastructure software for Oracle Domain Services Cluster.



(i) Note

Oracle Domain Services Cluster is not supported on Microsoft Windows.

Configure as an Oracle Extended Cluster: Select this option to extend an Oracle RAC cluster across two or more separate sites, each equipped with its own storage.

- In the Grid Plug and Play Information window, select SCAN option for the local cluster:
 - Select Create Local SCAN to create a SCAN locally. Enter the values in the Cluster Name, SCAN Name, and SCAN Port fields that are unique throughout your entire enterprise network.
 - 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 servers.
 - Select Use Shared GNS if you have configured a shared GNS on one cluster to provide name resolution for all clusters. Specify the complete path to the GNS client data file in the GNS Client Data field.
- Respond to the installation screens that appear in response to your configuration selection. Installation screens vary depending on the installation option you select.

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.

The following is a list of additional information about node IP addresses:

- For the local node only, the installer automatically fills in public and VIP fields. If your system uses vendor clusterware, then the installer 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 the installer 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 hostname command.

- Select storage option for Oracle Cluster Registry (OCR) and voting files:
 - Select Use Oracle Flex ASM for storage to store OCR and voting files on an Oracle ASM disk group.
 - Select Configure as ASM Client Cluster to store OCR and voting files on an Oracle ASM disk group configured on a storage server cluster. Specify the complete path to the ASM client data file in the ASM Client Data field.
- 10. In the Create GIMR Option window, select an option to create a Grid Infrastructure Management Repository for your Oracle Standalone Cluster installation:
 - Select Yes to create a GIMR by installing Oracle RAC software and creating GIMR container database after completing Oracle Grid Infrastructure installation.



- Select Use an Existing GIMR database to use a remote GIMR and specify a GIMR credentials file.
- Select No if you do not want to create a GIMR.

If you choose **Yes** on this window, then the Grid Infrastructure Management Repository Option window appears. Otherwise, the Create ASM Disk Group or ASM Client Storage Option window appears, based on your storage option. If you choose **No** on this window and you create a GIMR later, then create a separate Oracle ASM disk group for GIMR using Oracle ASMCA.

- 11. Choose whether you want to store the Grid Infrastructure Management Repository in a separate Oracle ASM disk group.
- 12. Provide the name and specifications for the Oracle ASM disk group.
 - In the Disk Group Name field, enter a name for the disk group, for example DATA.
 - Choose the Redundancy level for this disk group. Normal is the recommended option.
 - In the Add Disks section, choose the disks to add to this disk group.
 - If you want to use Oracle ASM Filter Driver (Oracle ASMFD) to manage your Oracle ASM disk devices, then select the option **Configure Oracle ASM Filter Driver**.
- 13. If you selected to use ASM client cluster in Step 9, then ASM Client Storage Option window appears. Select an Oracle ASM disk group from the storage server cluster to store the OCR and voting files.
- 14. Respond to the installation screens that appear in response to your configuration selection.
- 15. Provide information and run scripts as prompted by the installer.
 After the installation interview, you can click **Details** to see the log file.
- 16. After you have specified all the information needed for installation, the installer installs the software on each node. The installer then runs Oracle Net Configuration Assistant (NETCA), Oracle Private Interconnect Configuration Assistant, and Cluster Verification Utility (CVU). These programs run without user intervention.
- 17. During the installation, Oracle Automatic Storage Management Configuration Assistant (asmca) configures Oracle ASM for storage.

When you have verified that your Oracle Grid Infrastructure installation has completed successfully, you can either use Oracle Clusterware and Oracle ASM to maintain high availability for other applications, or you can install Oracle Database and Oracle RAC software.

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.

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 <code>setup.exe</code>, 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 9-1 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]
# node
             - Node's public host name
# vip
            - Node's virtual host name
# site-name
              - Node's assigned site
# Specify details of one node per line.
# Lines starting with '#' will be skipped.
# (1) vip is not required for Oracle Grid Infrastructure software only
# installs
# (2) vip should be specified as AUTO if Node Virtual host names are Dynamically
# (3) site-name should be specified only when configuring Oracle Grid Infrastructure with "Extended Cluster" option
# Examples:
# For installing GI software only on a cluster:
# ^^^^^^^^^^
# node1
```

node2



```
# For Standalone Cluster:
# ^^^^^^
# node1 node1-vip
# node2 node2-vip
# For Standalone Extended Cluster:
# ^^^^^^
# node1 node1-vip sitea
# node2 node2-vip siteb
```

Installing Software Binaries for Oracle Grid Infrastructure for a Cluster

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.

You can install Oracle Grid Infrastructure for a cluster software on multiple nodes at a time:

- Log in to Windows as the installation user for Oracle Grid Infrastructure, which must be a member of the Administrators users group.
- Create the Grid home directory.

For example:

```
md C:\app\21.0.0\grid
```

Ensure that the Grid home directory path you create is in compliance with the Oracle Optimal Flexible Architecture recommendations.

Download the Oracle Grid Infrastructure image files and copy the files to the Grid home.

For example:

```
cd C:\app\21.0.0\grid
unzip -q download_location\grid_home.zip
```

Unzip the installation image files only in the newly created Grid home directory.



(i) Note

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.

- Run the setup.exe command from the newly created Grid home directory, and select the Configuration Option as **Set Up Software Only**.
- 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.



- 6. When the software is configured, run the orainstRoot.bat script on all nodes, if prompted.
- Ensure that you have completed all storage and server preinstallation requirements.
- 8. Verify that all of the cluster nodes meet the installation requirements.

Use the command:

runcluvfy.bat stage -pre crsinst -n node_list

Configure the cluster using the Oracle Universal Installer (OUI) configuration wizard or configure the cluster using a response file.

Related Topics

Configuring the Software Binaries for Oracle Grid Infrastructure for a Cluster
 Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode.

Related Topics

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 using the Grid Setup Wizard utility (setup.exe).

Configuring the Software Binaries for Oracle Grid Infrastructure for a Cluster

Configure the software binaries by starting Oracle Grid Infrastructure configuration wizard in GUI mode.

After you have installed the Oracle Grid Infrastructure software on at least the local node, you can configure Oracle Grid Infrastructure on all the nodes in your cluster.

- Log in to one of the cluster nodes as the Oracle Installation user, and change directory to Grid_home.
- Start the Oracle Grid Infrastructure configuration wizard:

C:\Grid_home> setup.exe

The configuration script starts Oracle Universal Installer in Configuration Wizard mode.

3. Provide information as needed for configuration.

The configuration wizard mode validates the information and configures the installation on all cluster nodes.

4. Verify that the summary has the correct information for your cluster, and click **Install** to start configuration of the local node.

When you complete providing information, OUI shows you the Summary page, listing the information you have provided for the cluster.

When configuration of the local node is complete, OUI copies the Oracle Grid Infrastructure configuration file to other cluster member nodes.

Once all the root scripts are run, OUI checks the cluster configuration status, and starts other configuration tools as needed.



Related Topics

Oracle Clusterware Administration and Deployment Guide

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 using the Grid Setup Wizard utility (setup.exe).

Use this procedure to complete configuration after the software is installed or copied on nodes.

1. As the Oracle Installation user for Oracle Grid Infrastructure (for example, grid), start Oracle Universal Installer (OUI) in Oracle Grid Infrastructure Grid Setup Wizard mode from the Oracle Grid Infrastructure software-only home. Use the following syntax, where Grid_home is the Oracle Grid Infrastructure home, and filename is the response file name:

Grid_home\setup.exe [-debug] [-silent -responseFile filename]

For example:

C:\> cd Grid home

C:\> setup.exe -responseFile C:\app\21.0.0\grid\response\grid_setup.rsp

The configuration script starts OUI in Grid Setup Wizard mode. Each page shows the same user interface and performs the same validation checks that OUI normally does. However, instead of running an installation, the Grid Setup Wizard mode validates inputs and configures the installation on all cluster nodes.

When you complete inputs, OUI shows you the Summary page, listing all inputs you have provided for the cluster.

- 2. 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. OUI checks the cluster configuration status, and starts other configuration tools as needed.

(i) See Also

- Oracle Database Installation Guide for Microsoft Windows to configure and activate a software-only Oracle Grid Infrastructure installation for a standalone server
- Running Postinstallation Configuration Using a Response File

Setting Ping Targets for Network Checks

Receive notification about the status of public networks by setting the Ping_Targets parameter during the Oracle Grid Infrastructure installation.

In certain environments, for example, in a virtual machine, the network link status is not correctly returned when the network cable is disconnected. You can receive notification about



public network status in these environments by setting the Ping_Targets parameter during the Oracle Grid Infrastructure installation. You should use this parameter for addresses outside the cluster, like a switch or router.

Run the installer:

C:\..> setup.exe oracle_install_crs_Ping_Targets=Host1|IP1,Host2|IP2

For example:

C:\..> setup.exe oracle_install_crs_Ping_Targets=192.0.2.1,192.0.2.2

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.

Confirming Oracle Clusterware Function

After installation, use the crestl utility to verify Oracle Clusterware installation is installed and running correctly.

• Log in as a member of the Administrators group, and run the following command from the bin directory in the Grid home:

crsctl check cluster -all

Run the following command to verify the status of all the configured resources:

crsctl stat res -t

Example 9-2 Checking the Status of Oracle Clusterware

C:\..\bin\> crsctl check cluster -all

\$ 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



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



To manage Oracle ASM or Oracle Net, use the srvctl binary in the Oracle Grid Infrastructure home for a cluster (Grid home). If you have Oracle Real Application Clusters or Oracle Database installed, then you cannot use the srvctl binary in the database home to manage Oracle ASM or Oracle Net.

Understanding Offline Processes in Oracle Grid Infrastructure

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 postinstallation 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 the resource status:

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 will prompt you to activate the required offline resource.

Oracle Grid Infrastructure Postinstallation Tasks

Complete the postinstallation tasks after you have installed the Oracle Grid Infrastructure software.

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.

① 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
 Section and installation tasks
 - Certain postinstallation tasks are critical for your newly installed software.
- Recommended Postinstallation Tasks
 Oracle recommends that you complete these tasks as needed after installing Oracle Grid Infrastructure.
- Installing Grid Infrastructure Management Repository
 Install and configure Grid Infrastructure Management Repository (GIMR) into a separate
 Oracle home in this release of Oracle Grid Infrastructure.
- <u>Using Earlier Oracle Database Releases with Grid Infrastructure</u>
 Review the guidelines and restrictions for using earlier Oracle Database releases with Oracle Grid Infrastructure 21c installations.
- Modifying Oracle Clusterware Binaries After Installation
 After installation, if you must modify the software installed in your Grid home, then you must first stop the Oracle Clusterware stack.

Required Postinstallation Tasks

Certain postinstallation tasks are critical for your newly installed software.

patches for your Oracle software after you complete installation.



Backing up a voting file is no longer required.

<u>Downloading Release Update Patches</u>
 Download and install Release Updates (RU) and Release Update Revisions (RUR)



Configuring Exceptions for the Windows Firewall

To enable Oracle software to accept connection requests, you must configure the Windows Firewall by either opening up specific static transmission control protocol (TCP) ports in the firewall or by creating exceptions for specific executable files so they can receive connection requests on any ports they choose.

Downloading Release Update Patches

Download and install Release Updates (RU) and Release Update Revisions (RUR) patches for your Oracle software after you complete installation.

Starting with Oracle Database 18c, Oracle provides quarterly updates in the form of Release Updates (RU) and Release Update Revisions (RUR). Oracle no longer releases patch sets. For more information, see My Oracle Support Note 2285040.1.

Check the My Oracle Support website for required updates for your installation.

1. Use a web browser to view the My Oracle Support website:

https://support.oracle.com

2. Log in to My Oracle Support website.



Note

If you are not a My Oracle Support registered user, then click Register for My Oracle Support and register.

- On the main My Oracle Support page, click **Patches & Updates**.
- In the Patch Search region, select Product or Family (Advanced).
- 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.

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.

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

Configuring Exceptions for the Windows Firewall

To enable Oracle software to accept connection requests, you must configure the Windows Firewall by either opening up specific static transmission control protocol (TCP) ports in the firewall or by creating exceptions for specific executable files so they can receive connection requests on any ports they choose.

Review the list the Oracle Database executable files that listen on TCP ports on Windows, along with a brief description of the executable file. It is recommended to add these executable files (if in use and accepting connections from a remote, client computer) to the exceptions list for the Windows Firewall to ensure correct operation. In addition, if multiple Oracle homes are



in use, firewall exceptions may have to be created for the same executable file, for example, oracle.exe. multiple times, once for each Oracle home from which that executable file loads.

- About Configuring Exceptions for the Windows Firewall
 You must configure exceptions for the Windows Firewall to allow successful incoming connections to the Oracle software.
- <u>Firewall Exceptions for Oracle Database</u>
 For basic database operation and connectivity from remote clients, such as SQL*Plus,
 Oracle Call Interface (OCI), Open Database Connectivity (ODBC), and so on, you must
 add executable files to the Windows Firewall exception list.
- <u>Firewall Exceptions for Oracle Database Examples (or the Companion CD)</u>
 After installing the Oracle Database Companion CD, you must add executable files to the Windows Firewall exception list.
- Firewall Exceptions for Oracle Gateways
 If your Oracle database interacts with non-Oracle software through a gateway, then you must add the gateway executable file to the Windows Firewall exception list.
- Firewall Exceptions for Oracle Clusterware and Oracle ASM

 If you installed the Oracle Grid Infrastructure software on the nodes in your cluster, then you can enable the Windows Firewall only after adding certain executable files and ports to the Firewall exception list.
- <u>Firewall Exceptions for Oracle RAC Database</u>
 After installing the Oracle Real Application Clusters (Oracle RAC), you must add executable files to the Windows Firewall exception list.
- Firewall Exceptions for Other Oracle Products
 In additional to all the previously listed exceptions, if you use any of the Oracle software listed in, then you must create an exception for Windows Firewall for the associated executable file.
- <u>Troubleshooting Windows Firewall Exceptions</u>
 If you cannot establish certain connections even after granting exceptions to the executable files, then follow these steps to troubleshoot the installation.

About Configuring Exceptions for the Windows Firewall

You must configure exceptions for the Windows Firewall to allow successful incoming connections to the Oracle software.

If the Windows Firewall feature is enabled on one or more nodes in your cluster, then virtually all transmission control protocol (TCP) network ports are blocked to incoming connections. Any Oracle product that listens for incoming connections on a TCP port does not receive any of those connection requests, and the clients making those connections report errors unless you configure exceptions for the Windows Firewall.

You must configure exceptions for the Windows Firewall if your system meets all of the following conditions:

- Oracle server-side components are installed on a computer running a supported version of Microsoft Windows. The list of components includes the Oracle Database, Oracle Grid Infrastructure, Oracle Real Application Clusters (Oracle RAC), network listeners, or any web servers or services.
- The Windows computer accepts connections from other computers over the network. If no
 other computers connect to the Windows computer to access the Oracle software, then no
 post-installation configuration steps are required and the Oracle software functions as
 expected.



• The Windows computer is configured to run the Windows Firewall. If the Windows Firewall is not enabled, then no post-installation configuration steps are required.

Use one of the following methods to configure the firewall:

- Start the Windows Firewall application, select the Exceptions tab and then click either
 Add Program or Add Port to create exceptions for the Oracle software.
- From the command prompt, run the netsh firewall add... command.

A notification appears that a foreground application is attempting to listen on a port, and you are prompted to create an exception for that executable file. If you choose to create the exception in this way, the effect is the same as creating an exception for the executable file either through Control Panel or from the command line.

Firewall Exceptions for Oracle Database

For basic database operation and connectivity from remote clients, such as SQL*Plus, Oracle Call Interface (OCI), Open Database Connectivity (ODBC), and so on, you must add executable files to the Windows Firewall exception list.

The following executable files must be added to the Windows Firewall exception list:

- Oracle_home\bin\oracle.exe Oracle Database executable
- Oracle_home\bin\tnslsnr.exe Oracle Listener

If you use remote monitoring capabilities for your database, the following executable file must be added to the Windows Firewall exception list:

Oracle_home\jdk\bin\java.exe - Java Virtual Machine (JVM) for Oracle Enterprise Manager

Firewall Exceptions for Oracle Database Examples (or the Companion CD)

After installing the Oracle Database Companion CD, you must add executable files to the Windows Firewall exception list.

The following executable files must be added to the Windows Firewall exception list:

- Oracle_home\opmn\bin\opmn.exe Oracle Process Manager
- Oracle_home\jdk\bin\java.exe JVM

Firewall Exceptions for Oracle Gateways

If your Oracle database interacts with non-Oracle software through a gateway, then you must add the gateway executable file to the Windows Firewall exception list.

The following table lists the gateway executable files used to access non-Oracle software.

Table 10-1 Oracle Executables Used to Access Non-Oracle Software

Executable Name	Description
omtsreco.exe	Oracle Services for Microsoft Transaction Server
dg4sybs.exe	Oracle Database Gateway for Sybase
dg4tera.exe	Oracle Database Gateway for Teradata
dg4msql.exe	Oracle Database Gateway for SQL Server



Table 10-1 (Cont.) Oracle Executables Used to Access Non-Oracle Software

Executable Name	Description
dg4db2.exe	Oracle Database Gateway for Distributed Relational Database Architecture (DRDA)
pg4arv.exe	Oracle Database Gateway for Advanced Program to Program Communication (APPC)
pg4t4ic.exe	Oracle Database Gateway for APPC
dg4mqs.exe	Oracle Database Gateway for WebSphere MQ
dg4mqc.exe	Oracle Database Gateway for WebSphere MQ
dg4odbc.exe	Oracle Database Gateway for ODBC

Firewall Exceptions for Oracle Clusterware and Oracle ASM

If you installed the Oracle Grid Infrastructure software on the nodes in your cluster, then you can enable the Windows Firewall only *after* adding certain executable files and ports to the Firewall exception list.

The Firewall exception list must be updated on each node.

- *Grid_home*\bin\gpnpd.exe Grid Plug and Play daemon
- Grid_home\bin\oracle.exe Oracle Automatic Storage Management (Oracle ASM)
 executable file (if using Oracle ASM for storage)
- Grid_home\bin\racgvip.exe Virtual Internet Protocol Configuration Assistant
- Grid_home\bin\evmd.exe OracleEVMService
- Grid home\bin\crsd.exe OracleCRService
- Grid_home\bin\ocssd.exe OracleCSService
- Grid_home\bin\octssd.exe Cluster Time Synchronization Service daemon
- Grid_home\bin\mDNSResponder.exe multicast-domain name system (DNS) Responder Daemon
- Grid_home\bin\gipcd.exe Grid inter-process communication (IPC) daemon
- Grid_home\bin\gnsd.exe Grid Naming Service (GNS) daemon
- Grid_home\bin\ohasd.exe OracleOHService
- Grid_home\bin\TNSLSNR.EXE single client access name (SCAN) listener and local listener for Oracle RAC database and Oracle ASM
- Grid home\opmn\bin\ons.exe Oracle Notification Service (ONS)
- Grid_home\jdk\jre\bin\java.exe JVM
- Grid_home\bin\orangent.exe Oracle Agent
- Grid_home\bin\evmlogger.exe
- Grid home\bin\cssdagent.exe
- *Grid_home*\bin\cssdmonitor.exe
- Grid_home\bin\osysmond.exe System Monitor Service



- Grid_home\bin\orarootagent.exe Oracle Root Agent
- Grid_home\bin\ologgerd.exe Cluster Logger Service

Firewall Exceptions for Oracle RAC Database

After installing the Oracle Real Application Clusters (Oracle RAC), you must add executable files to the Windows Firewall exception list.

For the Oracle RAC database, the executable file that require exceptions are:

- Oracle home\bin\oracle.exe Oracle RAC database instance
- Oracle_home\jdk\bin\java.exe For the Oracle Enterprise Manager Database Console

In addition, the following ports should be added to the Windows Firewall exception list:

- Microsoft file sharing system management bus (SMB)
 - TCP ports from 135 through 139
- Direct-hosted SMB traffic without a network basic I/O system (NetBIOS)
 - port 445 (TCP)

Firewall Exceptions for Other Oracle Products

In additional to all the previously listed exceptions, if you use any of the Oracle software listed in, then you must create an exception for Windows Firewall for the associated executable file.

Table 10-2 Other Oracle Software Products Requiring Windows Firewall Exceptions

Oracle Software Product	Executable Name
Data Guard Manager	dgmgrl.exe
Oracle Internet Directory lightweight directory access protocol (LDAP) Server	oidldapd.exe
External Procedural Calls	extproc.exe

Troubleshooting Windows Firewall Exceptions

If you cannot establish certain connections even after granting exceptions to the executable files, then follow these steps to troubleshoot the installation.

- **1.** Examine Oracle configuration files (such as *.conf files), the Oracle key in the Windows registry, and network configuration files in %ORACLE_HOME%\network\admin.
- Grant an exception in the Windows Firewall to any executable listed in %ORACLE HOME%\network\admin\listener.ora in a PROGRAM= clause.
 - Each of these executables must be granted an exception in the Windows Firewall because a connection can be made through the TNS listener to that executable.
- 3. Examine Oracle trace files, log files, and other sources of diagnostic information for details on failed connection attempts.

Log and trace files on the database client computer may contain useful error codes or troubleshooting information for failed connection attempts. The Windows Firewall log file on the server may contain useful information as well.



4. If the preceding troubleshooting steps do not resolve a specific configuration issue on Windows, then provide the output from the following command to Oracle Support for diagnosis and problem resolution:

netsh firewall show state verbose=enable

Recommended Postinstallation Tasks

Oracle recommends that you complete these tasks as needed after installing Oracle Grid Infrastructure.

- About Installing Oracle Autonomous Health Framework
 Oracle Autonomous Health Framework is pre-installed with Oracle Database.
- Optimize Memory Usage for Programs
 The Windows operating system should be optimized for Memory Usage of 'Programs' instead of 'System Caching'.
- Create a Fast Recovery Area Disk Group
 You should create a separate disk group for the fast recovery area.
- Checking the SCAN Configuration
 The SCAN is a name that provides service access for clients to the cluster. 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.

About Installing Oracle Autonomous Health Framework

Oracle Autonomous Health Framework is pre-installed with Oracle Database.

Oracle Autonomous Health Framework performs proactive heath checks and collects diagnostics data for the Oracle software stack. Oracle Autonomous Health Framework updates are available on My Oracle Support and also shipped with Oracle Database Release Updates.

Oracle recommends that you update to the latest version of Oracle Autonomous Health Framework by either applying the latest Database Release Updates or by downloading and installing it from My Oracle Support Note 2550798.1:

https://support.oracle.com/epmos/faces/DocContentDisplay? id=2550798.1&parent=DOCUMENTATION&sourceId=USERGUIDE

Oracle Autonomous Health Framework includes the functionality from Oracle ORAchk, Oracle EXAchk, and Oracle Trace File Analyzer (TFA). Oracle Autonomous Health Framework extends health check coverage to the entire Oracle software stack, based on critical and reoccurring problems. Oracle Autonomous Health Framework proactively scans for known problems with Oracle products and deployments, including the following:

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



- Oracle SuperCluster
- Oracle Zero Data Loss Recovery Appliance

Related Topics

Oracle Autonomous Health Framework Checks and Diagnostics User's Guide

Optimize Memory Usage for Programs

The Windows operating system should be optimized for Memory Usage of 'Programs' instead of 'System Caching'.

- 1. From the Start Menu, select Control Panel, then System.
- 2. In the System Properties window, click the **Advanced** tab.
- 3. In the Performance section, click **Settings**.
- 4. In the Performance Options window, click the Advanced tab.
- 5. In the Memory Usage section, ensure Programs is selected.

Create a Fast Recovery Area Disk Group

You should create a separate disk group for the fast recovery area.

During installation of Oracle Grid Infrastructure, if you select Oracle ASM for storage, a single disk group is created to store the Oracle Clusterware files. If you plan to create a single-instance database, an Oracle RAC database, or an Oracle RAC One Node database, then this disk group can also be used to store the data files for the database. However, Oracle recommends that you create a separate disk group for the fast recovery area.

- 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.
- Creating the Fast Recovery Area Disk Group
 You can use ASMCA to create an Oracle ASM disk group for the fast recovery area.

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.

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. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

When you enable the fast recovery area in the database initialization parameter file, all RMAN backups, archive logs, control file automatic backups, and database copies are written to the fast recovery area. RMAN automatically manages files in the fast recovery area by deleting obsolete backups and archive files that are no longer required for recovery.

To use a fast recovery area in Oracle RAC, you must place it on an Oracle ASM disk group, a cluster file system, or on a shared directory that is configured through Direct network file system (NFS) for each Oracle RAC instance. In other words, the fast recovery area must be shared among all of the instances of an Oracle RAC database. Oracle Clusterware files and Oracle Database files can be placed on the same disk group as fast recovery area files.



However, Oracle recommends that you create a separate fast recovery area disk group to reduce storage device contention.

The fast recovery area is enabled by setting the parameter DB RECOVERY FILE DEST to the same value on all instances. The size of the fast recovery area is set with the parameter 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 should be 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 one fast recovery area disk group on disks with 150 gigabyte (GB) of storage, shared by three different databases. You can set the size of the fast recovery area for each database depending on the importance of each database. For example, if test1 is your least important database, you might set DB_RECOVERY_FILE_DEST_SIZE to 30 GB. For the products database, which is of greater importance, you might set DB RECOVERY FILE DEST SIZE to 50 GB. For the orders, which has the greatest importance, you might set DB RECOVERY FILE DEST SIZE to 70 GB.

(i) See Also

Oracle Automatic Storage Management Administrator's Guide for information on how to create a disk group for data and a disk group for the fast recovery area

Creating the Fast Recovery Area Disk Group

You can use ASMCA to create an Oracle ASM disk group for the fast recovery area.

Navigate to the bin directory in the Grid home and start Oracle ASM Configuration Assistant (ASMCA).

For example:

 $C:\> cd C:\app\20.0.0\grid\bin$ C:\> asmca

ASMCA opens at the Disk Groups tab.

- Click **Disk Groups** in the left panel to open the Disk Groups tab.
- Click **Create** to create a new disk group.

The Create Disk Groups window opens.

- In the Create Disk Groups window, enter the following information, then click **OK**:
 - a. In the Disk Group Name field, enter a descriptive name for the fast recovery area disk group. For example: FRA.
 - b. In the Redundancy section, select the level of redundancy you want to use. For example: Normal.
 - c. In the Select Member Disks field, select eligible disks to be added to the fast recovery area, and click **OK**.



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 SCAN is a name that provides service access for clients to the cluster. 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.

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 run a particular database instance. Clients can continue to access the cluster in the same way as with earlier releases, but Oracle recommends that clients accessing the cluster use the SCAN.

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

Confirm that the DNS is correctly associating the SCAN with the specified addresses.

cluvfy comp scan

Example 10-1 Using CLUVFY to Confirm DNS is Correctly Associating the SCAN Addresses

This example shows the output from the cluvfy comp scan command for a cluster node named node1.example.com.

C:\> cluvfy comp scan Verifying scan

Checking Single Client Access Name (SCAN)...

Checking TCP connectivity to SCAN Listeners...
TCP connectivity to SCAN Listeners exists on all cluster nodes

Checking name resolution setup for "node1.example.com"...

Verification of SCAN VIP and Listener setup passed

Verification of scan was successful.

Installing Grid Infrastructure Management Repository

Install and configure Grid Infrastructure Management Repository (GIMR) into a separate Oracle home in this release of Oracle Grid Infrastructure.

- About Grid Infrastructure Management Repository Installation
 You can install Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster.
- Creating the Grid Infrastructure Management Repository Disk Group
 Create an Oracle Automatic Storage Management (Oracle ASM) disk group to host the GIMR for a cluster.



- Installing Grid Infrastructure Management Repository for a New Cluster Complete this procedure to install Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster with local GIMR.
- Adding Grid Infrastructure Management Repository to an Existing Cluster Complete this procedure to add a Grid Infrastructure Management Repository to an existing cluster.

About Grid Infrastructure Management Repository Installation

You can install Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster.

Starting with Oracle Grid Infrastructure 21c, the Grid Infrastructure Management Repository will be configured in a separate Oracle home, instead of Oracle Grid Infrastructure home (Grid home). This option is available if you choose to configure GIMR during a fresh installation or you add a GIMR to an existing Oracle Grid Infrastructure deployment.



Note

Installing GIMR is optional for Oracle Standalone Clusters.

Install and configure Grid Infrastructure Management Repository in a new Oracle home after installing and configuring Oracle Grid infrastructure. After you install Oracle Grid Infrastructure and choose to install GIMR, you need to download and install Oracle Database software of the same version as Oracle Grid Infrastructure.

If you choose to install GIMR during Oracle Grid Infrastructure installation, then you are prompted to create a new Oracle ASM disk group for GIMR and your decision is recorded in Oracle Cluster Registry (OCR). Cluster Verification Utility (CVU) creates a log to remind you to install GIMR. However, if you do not choose to install GIMR during Oracle Grid Infrastructure installation, then CVU does not remind you to create a GIMR.



(i) Note

You can choose to use the same disk as used for Oracle Grid Infrastructure (OCR) or you can use a new disk for GIMR.

Creating the Grid Infrastructure Management Repository Disk Group

Create an Oracle Automatic Storage Management (Oracle ASM) disk group to host the GIMR for a cluster.

If you have not created an Oracle ASM disk group for GIMR during Oracle Grid Infrastructure installation, then follow these steps to create a new Oracle ASM disk group using Oracle ASMCA.

Go to the Oracle Grid Infrastructure home bin directory, and start Oracle ASM Configuration Assistant (ASMCA).

 $C:\> cd C:\app\21.0.0\grid\bin$

C:\> asmca



ASMCA home window opens.

- Click **Disk Groups** in the left panel to open the Disk Groups tab.
- Click **Create** to create a new disk group.

The Create Disk Groups window opens.

Provide configuration information for the GIMR disk group as prompted:

In the Disk Group Name field, enter a descriptive name for the GIMR disk group. For example: MGMT.

In the Redundancy section, select the level of redundancy you want to use. For example: External

In the Select Member Disks field, select eligible disks you want to add to the MGMT disk group, and click OK.



(i) Note

Determine the minimum storage space that you require for the GIMR and make sure the disk group meets the minimum storage space requirements for your GIMR.

When the GIMR disk group creation is complete, click **Exit** and click **Yes** to confirm closing the ASMCA application.

Installing Grid Infrastructure Management Repository for a New Cluster

Complete this procedure to install Grid Infrastructure Management Repository (GIMR) for Oracle Standalone Cluster with local GIMR.

If you want to place GIMR on an Oracle ASM disk group, then create an Oracle ASM disk group for GIMR either during Oracle Grid Infrastructure installation or using Oracle ASMCA before installing the GIMR.

As the grid user, create GIMR home directory on all cluster nodes:

 $C:\$ md $C:\$ pp\21.0.0\gimr

Download the Oracle Database image files and extract the files into the GIMR home:

 $C:\> cd C:\app\21.0.0\gimr$ C:\> unzip -q download_location\db_home.zip

db_home.zip is the name of the Oracle Database image zip file.



(i) Note

- You must extract the zip image software into the directory where you want your GIMR home to be located.
- Download and copy the Oracle Database image files to the local node only. During installation, the software is copied and installed on all other nodes in the cluster.



From the GIMR home directory, start the Oracle Database software installation:

 $C:\> cd C:\app\21.0.0\gimr$ C:\> setup.exe



Note

Run the setup.exe command from the GIMR home directory only. Do not run setup.exe command that resides at %GIMR_HOME%\oui\bin\ or any other location to install Oracle Real Application Clusters.

- In the Select Configuration Option screen, select the Setup Software Only option. Click Next.
- In the Select Database Installation Option screen, select the **Oracle Real Application** Clusters Database Installation option. Click Next.
- In the Node Selection screen, select all the nodes to install Oracle RAC or Oracle RAC One Node software. Click Next.
- In the Select Database Edition screen, select the Enterprise Edition option. Click Next.
- In the Specify Install Location screen, provide the location of Oracle base for Oracle RAC software installation.
 - The Oracle base directory must be different from the GIMR home directory. Click **Next**.
- The Perform Prerequisite Checks screen displays the results of the prerequisites checks. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the Check Again button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click Next.
- 10. Review the contents of the Summary screen, and then click Install.
 - The installer displays a progress indicator enabling you to monitor the installation process.
- 11. Create GIMR container database using the following command:

C:\> C:\app\21.0.0\gimr\bin\mgmtca createGIMRContainer [-storageDiskLocation disk_location]

Adding Grid Infrastructure Management Repository to an Existing Cluster

Complete this procedure to add a Grid Infrastructure Management Repository to an existing cluster.

If you want to place GIMR on an Oracle ASM disk group, then create an Oracle ASM disk group for GIMR either during Oracle Grid Infrastructure installation or using Oracle ASMCA before installing the GIMR.

As the grid user, create GIMR home directory on all cluster nodes:

 $C:\$ md $C:\$ pp\21.0.0\gimr



Download the Oracle Database image files and extract the files into the GIMR home. For example:

C:\> cd C:\app\21.0.0\gimr C:\> unzip -q download_location\db_home.zip

db_home.zip is the name of the Oracle Database image zip file.

Note

- You must extract the zip image software into the directory where you want your GIMR home to be located.
- Download and copy the Oracle Database image files to the local node only.
 During installation, the software is copied and installed on all other nodes in the cluster.
- 3. From the GIMR home directory, start the Oracle Database software installation:

C:\> cd C:\app\21.0.0\gimr C:\> setup.exe

Note

Run the setup.exe command from the GIMR home directory only. Do not run setup.exe command that resides at %GIMR_HOME%\oui\bin\, or any other location, to install Oracle Real Application Clusters.

- In the Select Configuration Option screen, select the Setup Software Only option. Click Next.
- In the Select Database Installation Option screen, select the Oracle Real Application Clusters Database Installation option. Click Next.
- 6. In the Node Selection screen, select all the nodes to install Oracle RAC or Oracle RAC One Node software. Click **Next**.
- 7. In the Select Database Edition screen, select the Enterprise Edition option. Click Next.
- In the Specify Install Location screen, provide the location of Oracle base for Oracle RAC software installation.

The Oracle base directory must be different from the GIMR home directory. Click Next.

- 9. The Perform Prerequisite Checks screen displays the results of the prerequisites checks. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the Check Again button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click Next.
- 10. Review the contents of the Summary screen and then click **Install**.

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



11. Create GIMR container database using the following command:

C:\> C:\app\21.0.0\gimr\bin\mgmtca createGIMRContainer [-storageDiskLocation disk_location]

Using Earlier Oracle Database Releases with Grid Infrastructure

Review the guidelines and restrictions for using earlier Oracle Database releases with Oracle Grid Infrastructure 21c installations.

- General Restrictions for Using Earlier Oracle Database Releases
 Restrictions for using earlier Oracle Database releases with Oracle Grid Infrastructure 21c.
- Using the Correct LSNRCTL Commands
 Do not attempt to use the Isnrctl programs from Oracle home locations for earlier releases because they cannot be used with the new release.

General Restrictions for Using Earlier Oracle Database Releases

Restrictions for using earlier Oracle Database releases with Oracle Grid Infrastructure 21c.

You can use Oracle Database 12c Release 1 and Release 2, Oracle Database 18c, Oracle Database 19c, and Oracle Database 21c with Oracle Grid Infrastructure 21c.

Do not use the versions of srvctl, Isnrctl, 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.

When you install Oracle databases on a Flex ASM cluster, you must set the Oracle ASM cardinality to All.

Using the Correct LSNRCTL Commands

Do not attempt to use the lsnrctl programs from Oracle home locations for earlier releases because they cannot be used with the new release.

 Use the Listener Control utility, Isnrctl, located in the Oracle Grid Infrastructure 21c home to administer local and SCAN listeners for Oracle Clusterware and Oracle ASM.

Modifying Oracle Clusterware Binaries After Installation

After installation, if you must modify the software installed in your Grid home, then you must first stop the Oracle Clusterware stack.

For example, to apply a one-off patch, or modify any of the dynamic-link libraries (DLLs) used by Oracle Clusterware or Oracle ASM, you must follow these steps to stop and restart Oracle Clusterware.



△ Caution

To put the changes you make to the Oracle Grid Infrastructure home into effect, you must shut down all executable files that run in the Grid home directory and then restart them. In addition, shut down any applications that use Oracle shared libraries or DLL files in the Grid home.

- 1. Log in using a member of the Administrators group and go to the directory *Grid_home*\bin, where *Grid_home* is the path to the Oracle Grid Infrastructure home.
- 2. Shut down Oracle Clusterware using the following command:

C:\..\bin> crsctl stop crs -f

- **3.** After Oracle Clusterware is completely shut down, perform the updates to the software installed in the Grid home.
- 4. Use the following command to restart Oracle Clusterware:

C:\..\bin> crsctl start crs

5. Repeat steps 1 through 4 on each cluster member node.

(i) Note

Do not delete directories in the Grid home. For example, do not delete the directory $Grid_home/OP$ atch. 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 $Grid_home/OP$ atch".

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.

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

- Understanding Out-of-Place and Rolling Upgrades
 You can use rolling upgrade or out-of-place upgrades to upgrade your Oracle Grid Infrastructure software.
- About Oracle Grid Infrastructure Upgrade and Downgrade
 There are different methods you can use to upgrade the Oracle Grid Infrastructure software.
- Options for Oracle Grid Infrastructure Upgrades
 Review these upgrade options when you upgrade to Oracle Grid Infrastructure 21c.
- Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades
 Review the restrictions and changes for upgrades to Oracle Grid Infrastructure installations.
- 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.
- <u>Understanding Rolling Upgrades Using Batches</u>
 Instead of shutting down all nodes when applying patches, you can shut down some nodes while other nodes remaining running.
- <u>Performing Rolling Upgrades of Oracle Grid Infrastructure</u>
 Review this information to perform rolling upgrade of Oracle Grid Infrastructure.
- Upgrading Grid Infrastructure Management Repository from an Earlier Release
 Complete this procedure to upgrade Grid Infrastructure Management Repository (GIMR) from an earlier release.
- Applying Patches to Oracle Grid Infrastructure
 After you have upgraded to Oracle Grid Infrastructure 21c, you can install individual software patches by downloading them from My Oracle Support.
- Applying Patches to Grid Infrastructure Management Repository
 You can patch Grid Infrastructure Management Repository by installing in-place or out-of-place patches using the OPatch utility.
- <u>Updating Oracle Enterprise Manager Cloud Control Target Parameters</u>
 After upgrading Oracle Grid Infrastructure, upgrade the Enterprise Manager Cloud Control target.
- Checking Cluster Health Monitor Repository Size After Upgrading
 If you are upgrading Oracle Grid Infrastructure from a prior release using IPD/OS to the current release, then review the Cluster Health Monitor repository size (the CHM repository).



- Downgrading Oracle Clusterware to an Earlier Release After a successful or a failed upgrade, you can downgrade Oracle Clusterware to the previous release.
- Completing Failed or Interrupted Installations and Upgrades If Oracle Universal Installer (OUI) exits on the node from which you started the installation or upgrade (the first node), or the node reboots before you confirm that the gridconfig.bat script was run on all cluster nodes, then the upgrade or installation remains incomplete.

Related Topics

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

Understanding Out-of-Place and Rolling Upgrades

You can use rolling upgrade or out-of-place upgrades to upgrade your Oracle Grid Infrastructure software.

If you have an existing Oracle Grid Infrastructure installation, then you upgrade your existing cluster by performing an out-of-place upgrade. An in-place upgrade of Oracle Grid Infrastructure is not supported. 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 also perform the upgrade in a rolling manner, which means there is always at least one cluster node operational during the upgrade.

Rolling Upgrades

You can upgrade Oracle Grid Infrastructure by upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster, which is called performing a **rolling** upgrade. Rolling upgrades avoid downtime and ensure continuous availability while the software is upgraded to a new release.



(i) Note

In contrast with releases prior to Oracle Clusterware 11g Release 2, Oracle Universal Installer (OUI) always performs rolling upgrades, even if you select all nodes for the upgrade.

Out-of-Place Upgrades

During an out-of-place upgrade, the installer installs the newer release in a separate Grid home. Both the old and new releases of Oracle Grid Infrastructure exist on each cluster member node, but only one release is active. By contrast, an in-place upgrade overwrites the software in the current Oracle Grid Infrastructure home.

To perform an out-of-place upgrade, you must create new Oracle Grid Infrastructure homes on each node. Then you can perform an out-of-place rolling upgrade, so that some nodes run Oracle Grid Infrastructure from the original Grid home, and other nodes run Oracle Grid Infrastructure from the new Grid home.

Related Topics

Performing Rolling Upgrades of Oracle Grid Infrastructure Review this information to perform rolling upgrade of Oracle Grid Infrastructure.



About Oracle Grid Infrastructure Upgrade and Downgrade

There are different methods you can use to upgrade the Oracle Grid Infrastructure software.

You can upgrade Oracle Grid Infrastructure in any of the following ways:

- Rolling Upgrade: Involves upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster
- Non-Rolling Upgrade: Involves bringing down all the nodes except one. A complete cluster
 outage occurs while the root script stops the old Oracle Clusterware stack and starts the
 new Oracle Clusterware stack on the node where you initiate the upgrade. After upgrade is
 completed, the new Oracle Clusterware is started on all the nodes.

Note that some services are disabled when one or more nodes are in the process of being upgraded. All upgrades are **out-of-place upgrades**, meaning that the software binaries are placed in a different Grid home from the Grid home used for the prior release.

You can downgrade from Oracle Grid Infrastructure 21c to Oracle Grid Infrastructue 19c, Oracle Grid Infrastructure 18c, Oracle Grid Infrastructure 12c Release 2 (12.2), and Oracle Grid Infrastructure 12c Release 1 (12.1). 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 ASMCA. In addition to running ASMCA using the graphical user interface, you can run ASMCA in non-interactive (silent) mode.

(i) Note

- If you are currently using OCFS for Windows as your shared storage, then you
 must migrate to using Oracle ASM during the upgrade of Oracle Database and
 Oracle Grid Infrastructure.
- You must complete an upgrade before attempting to use cluster backup files. You cannot use backups for a cluster that has not completed the upgrade.

See Also

Oracle Automatic Storage Management Administrator's Guide for additional information about upgrading existing Oracle ASM installations

Options for Oracle Grid Infrastructure Upgrades

Review these upgrade options when you upgrade to Oracle Grid Infrastructure 21c.

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

- Oracle Grid Infrastructure upgrade from 12c Release 2 (12.2) to Oracle Grid Infrastructure 21c.
- Oracle Grid Infrastructure upgrade from 18c release to Oracle Grid Infrastructure 21c.



Oracle Grid Infrastructure upgrade from 19c release to Oracle Grid Infrastructure 21c.

Upgrade options from Oracle Grid Infrastructure 12c Release 2 (12.2), Oracle Grid Infrastructure 18c, and Oracle Grid Infrastructure 19c to Oracle Grid Infrastructure 21c include the following:

- Oracle Grid Infrastructure rolling upgrade which involves upgrading individual nodes without stopping Oracle Grid Infrastructure on other nodes in the cluster
- Oracle Grid Infrastructure non-rolling upgrade by bringing the cluster down and upgrading the complete cluster

(i) Note

You can use either Oracle ASM or a shared file system to store OCR and voting
files on Oracle Standalone Cluster deployments. If storage for OCR and voting
files is other than Oracle ASM on other cluster types, then you need to migrate
OCR and voting files to Oracle ASM before upgrading to Oracle Grid Infrastructure
21c.

Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades

Review the restrictions and changes for upgrades to Oracle Grid Infrastructure installations.

- Oracle Grid Infrastructure upgrades are always out-of-place upgrades. You cannot perform an in-place upgrade of Oracle Grid Infrastructure to existing homes.
- You must use an Administrator user to perform the Oracle Grid Infrastructure 21c upgrade.
- Oracle ASM and Oracle Clusterware both run in the Oracle Grid Infrastructure home.
- When you upgrade to Oracle Grid Infrastructure 21c, you upgrade to an Oracle Flex Cluster configuration.
- Do not delete directories in the Grid home. For example, do not delete *Grid_home*\OPatch. If you delete the directory, then the Oracle Installation User for Oracle Grid Infrastructure cannot use Opatch to patch the Grid home, and Opatch displays the error "checkdir error: cannot create *Grid_home*\OPatch"
- To upgrade existing Oracle Grid Infrastructure installations to Oracle Grid Infrastructure 21c, you must first verify if you need to apply any mandatory patches for upgrade to succeed. You can use CVU to perform this check.
- During a major release upgrade to Oracle Clusterware 21c, the software in the Grid home for Oracle Grid Infrastructure 21c is not fully functional until the upgrade is completed.
 Running the Server Control Utility (SRVCTL), crsctl, and other commands from the 21c Grid home is not supported until the upgrade is complete across all nodes.

To manage databases using earlier releases of Oracle Database software during the Oracle Grid Infrastructure upgrade, use SRVCTL from the existing database homes.

Storage Restrictions Related to Oracle Grid Infrastructure Upgrades

• If the Oracle Cluster Registry (OCR) and voting file locations for your current installation are on raw devices or shared file systems, then you must migrate them to Oracle ASM disk groups before upgrading to Oracle Grid Infrastructure 21c.



Restrictions Related to Upgrading Shared Grid Homes

- You can perform upgrades on a shared Oracle Clusterware home.
- If the existing Oracle Clusterware home is a shared home, then you can use a non-shared home for the Oracle Grid Infrastructure for a cluster home for Oracle Clusterware and Oracle ASM 21c.

Single-Instance Oracle ASM Upgrade Restrictions

- During Oracle Grid Infrastructure installation or upgrade, if there is a single instance Oracle
 ASM release on the local node, then it is converted to an Oracle Flex ASM 21c installation,
 and Oracle ASM runs in the Oracle Grid Infrastructure home on all nodes.
- If a single instance (non-clustered) Oracle ASM installation is on a remote node, which is a
 node other than the local node (the node on which the Oracle Grid Infrastructure
 installation or upgrade is being performed), then it will remain 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.

Related Topics

- <u>Using CVU to Validate Readiness for Oracle Clusterware Upgrades</u>
 Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.
- <u>Example of Verifying System Upgrade Readiness for Grid Infrastructure</u>
 You can use the runcluvfy.bat command to check your system before upgrading.
- About the CVU Grid Upgrade Validation Command Options
 You can use the Cluster Verification Utility (CVU) to validate your system readiness before upgrading.

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.
- Tasks to Complete Before Upgrading Oracle Grid Infrastructure Review the tasks before upgrading Oracle Clusterware.
- Create an Oracle ASM Password File
 In certain situations, you must create an Oracle ASM password file before you can upgrade to Oracle Grid Infrastructure 21c.
- Running the Oracle ORAchk Upgrade Readiness Assessment
 Download and run the Oracle ORAchk Upgrade Readiness Assessment before upgrading Oracle Grid Infrastructure.



<u>Using CVU to Validate Readiness for Oracle Clusterware Upgrades</u>
 Oracle recommends that you use Cluster Verification Utility (CVU) to help to ensure that your upgrade is successful.

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.

If you intend to install or upgrade Oracle RAC, then you must first complete the upgrade to Oracle Grid Infrastructure 21c on all cluster nodes before you install the Oracle Database 21c release of Oracle RAC.

Check	Task		
Review the <i>Oracle Database Upgrade Guide</i> for deprecation and desupport information that may affect upgrade planning.	Oracle Database Upgrade Guide		
Verify the current operating system is supported for the new release	Confirm that you are using a supported operating system, patch release, and all required operating system packages for the new Oracle Grid Infrastructure installation.		
Patch set (recommended)	Install the latest patch set release for your existing installation. Upgrades from some earlier patch set releases might not be supported.		
Move any files stored on OCFS for Windows or raw devices to a supported storage mechanism	Migrate OCR files from raw devices to Oracle ASM. Direct use of raw devices is not supported. OCFS on Windows is no longer supported. Confirm Oracle Cluster Registry (OCR) file integrity. If this check fails, then repair the OCRs before proceeding.		
Install user account	Confirm that the installation user you plan to use is the same as the installation user that was used to perform the installation you want to upgrade.		
Oracle Home user	If the Oracle home from which you run the database upgrade uses a Windows Domain User as the Oracle Home User, then the Oracle Home User on the target version must use the same Windows Domain User.		
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	Confirm that the Oracle Automatic Storage Management (Oracle ASM) instances use standard Oracle ASM instance names		
	The default ASM SID for a single-instance database is +ASM, and the default SID for ASM on Oracle Real Application Clusters nodes is +ASM <i>node#</i> , where <i>node#</i> is the node number.		
	The nondefault Oracle ASM instance names are not supported. If you have nondefault Oracle ASM instance names, then before you upgrade your cluster, use your existing release srvctl to remove individual Oracle ASM instances with nondefault names and add Oracle ASM instances with default names.		



Check	Task	
Network Addresses for standard Oracle Grid Infrastructure	For standard Oracle Grid Infrastructure installations, confirm the following network configuration:	
deployments	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 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.*.*). 	
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 that will be performing the upgrade, unset the environment variables %ORACLE_HOME% and %ORACLE_SID% as follows:	
	C:\> set ORACLE_BASE=	
	C:\> set ORACLE_HOME=	
	C:\> set ORACLE_SID=	
	Check that the ORA_CRS_HOME environment variable is not set. Do	
	not use ORA_CRS_HOME as an environment variable, except under explicit direction from Oracle Support.	
RACcheck Upgrade Readiness Assessment	Download and run the RACcheck Upgrade Readiness Assessment to obtain automated upgrade-specific health check for upgrades to Oracle Grid Infrastructure. See My Oracle Support note 1457357.1, which is available at the following URL:	
	https://support.oracle.com/rs?type=doc&id=1457357.1	
Back Up the Oracle Software Before Upgrades	Before you make any changes to the Oracle software, Oracle recommends that you create a backup of the Oracle software and databases.	
	For Oracle software running on Windows operating systems, you must also take a backup of the Windows registry.	
	Without a registry backup, you will not be able to restore the Oracle software to a working state if the upgrade of Oracle Grid Infrastructure or Oracle Database fails and you want to revert to the previous software installation.	

Related Topics

Oracle Database Upgrade Guide

Tasks to Complete Before Upgrading Oracle Grid Infrastructure

Review the tasks before upgrading Oracle Clusterware.

Before upgrading, you must unset the following environment variables:

- ORACLE_BASE
- ORACLE_HOME
- ORACLE_SID



- ORA NLS10
- TNS ADMIN
- ORA_CRS_HOME

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.

- 1. For each node, use the Cluster Verification Utility (CVU) to assist you with system checks in preparation for patching or upgrading.
 - You can run CVU before starting the upgrade, however, the installer runs the appropriate CVU checks automatically, and prompts you to fix problems before proceeding with the upgrade.
- 2. Ensure that you have information you will 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 Grid home location
 - SCAN name and addresses, and other network addresses
 - Privileged user operating system groups
 - Local Administrator user access, or access as the user who performed the previous Oracle Clusterware installation
- 3. For the installation user running the installation, if you have environment variables set for the existing installation, then unset the environment variables %ORACLE_HOME% and %ORACLE_SID%, because these environment variables are used during upgrade. For example, as the grid user, run the following commands on the local node:

```
C:\> set ORACLE_HOME=
C:\> set ORACLE_BASE=
C:\> set ORACLE_SID=
```

- 4. 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.
- 5. If you have an existing installation on your system, and you are using the same user account to upgrade 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.
- Check to ensure that the user profile for the Oracle Installation User does not set any of these environment variables.

Related Topics

- Configuring Users, Groups and Environments for Oracle Grid Infrastructure and Oracle RAC
- Configuring Networks for Oracle Grid Infrastructure and Oracle RAC



Create an Oracle ASM Password File

In certain situations, you must create an Oracle ASM password file before you can upgrade to Oracle Grid Infrastructure 21c.

On Windows platforms, if you are upgrading Oracle Grid Infrastructure from release 11.2.0.4 to release 12.1.0.2, and then upgrade to Oracle Grid Infrastructure 21c, you must create an Oracle ASM password file before starting the upgrade to Oracle Grid Infrastructure 21c.

However, if you upgrade from Oracle Grid Infrastructure 12.1.0.2 directly to Oracle Grid Infrastructure 21c, then this task is not required.

If you do not complete this task before starting the upgrade process, then you will get a failed check during installation. If you choose to ignore the error and continue the upgrade, then the rooters script fails to complete and generates the following error:

2016-06-28 20:20:33: Command output: CLSRSC-661: The Oracle ASM password file does not exist at location C:\app\12.1.0\grid\database\PWD+ASM.ora.

After upgrading Oracle Grid Infrastructure from release 11.2.0.4 to release 12.1.0.2, complete the following steps to resolve this error before continuing with the upgrade to Oracle Grid Infrastructure 21c:

Start the ASMCMD utility from the upgraded Oracle Grid Infrastructure home directory.

Grid_home_12102\bin\asmcmd.bat

2. Set the Oracle ASM disk group compatibility to 12.1.0.0.0 or higher for the disk group that stores the Oracle Cluster Registry (OCR).

ASMCMD> setattr -G disk_group_name compatible.asm 12.1.0.0.0

Create an Oracle ASM password file.

ASMCMD> pwcreate --asm +disk_group_name/orapwASM sys_password

4. Verify the password file was created.

ASMCMD> pwget --asm +disk_group_name/orapwasm

Running the Oracle ORAchk Upgrade Readiness Assessment

Download and run the Oracle ORAchk Upgrade Readiness Assessment before upgrading Oracle Grid Infrastructure.

Oracle ORAchk is an Oracle RAC configuration audit tool. Oracle ORAchk Upgrade Readiness Assessment can be used to obtain an automated upgrade-specific health check for upgrades to Oracle Grid Infrastructure 12.2, 18c, 19c, and 21c. You can run the Oracle ORAchk Upgrade Readiness Assessment tool and automate many of the manual pre-upgrade and post-upgrade checks.



Oracle recommends that you download and run the latest version of Oracle ORAchk from My Oracle Support. For information about downloading, configuring, and running Oracle ORAchk, refer to My Oracle Support note 1457357.1.

Related Topics

- Oracle ORAchk and EXAchk User's Guide
- 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 Grid Upgrade Validation Command Options
 You can use the Cluster Verification Utility (CVU) to validate your system readiness before upgrading.
- <u>Example of Verifying System Upgrade Readiness for Grid Infrastructure</u>
 You can use the runcluvfy.bat command to check your system before upgrading.

Related Topics

Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades
 Review the restrictions and changes for upgrades to Oracle Grid Infrastructure installations.

About the CVU Grid Upgrade Validation Command Options

You can use the Cluster Verification Utility (CVU) to validate your system readiness before upgrading.

You can run upgrade validations in one of two ways:

- Run the installer, and allow the CVU validation built into the installer to perform system checks
- Run the CVU manual script cluvfy.bat script to perform system checks

To use the installer to perform pre-install checks, run the installation as you normally would. The installer starts CVU, and performs system checks as part of the installation process. Selecting the installer 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.bat command-line script for CVU, navigate to the new Grid home where you extracted the image files for upgrade, that contains the runcluvfy.bat script, and run the following command to check the readiness of your Oracle Clusterware installation for upgrades:

runcluvfy.bat stage -pre crsinst -upgrade

Running runcluvfy.bat 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 runcluvfy command uses the following syntax, where variable content is indicated by italics:

runcluvfy.bat stage -pre crsinst -upgrade [-rolling] -src_crshome src_Gridhome -dest_crshome dest_Gridhome -dest_version dest_release [-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 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: 12.1.0.1.0.

-verbose

Use the -verbose option to produce detailed output of individual checks.

Related Topics

Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades
 Review the restrictions and changes for upgrades to Oracle Grid Infrastructure installations.

Example of Verifying System Upgrade Readiness for Grid Infrastructure

You can use the runcluvfy.bat command to check your system before upgrading.

Verify that the permissions required for installing Oracle Grid Infrastructure have been configured by running the following command:

C:\app\21.0.0\grid> runcluvfy.bat stage -pre crsinst -upgrade -rolling -src_crshome C:\app\19.0.0\grid -dest_crshome C:\app\21.0.0\grid -dest version 21.0.0.0.0 -verbose

Related Topics

Restrictions and Guidelines for Oracle Grid Infrastructure Upgrades
 Review the restrictions and changes for upgrades to Oracle Grid Infrastructure installations.

Understanding Rolling Upgrades Using Batches

Instead of shutting down all nodes when applying patches, you can shut down some nodes while other nodes remaining running.



When you upgrade Oracle Grid Infrastructure, 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 upgrade process on each node, the upgrade scripts shut down the database instances and then start the instances again.

When performing the upgrade, you can divide the nodes into groups, or batches, and start upgrades of these node batches. Between batches, you can move services from nodes running the earlier release to the upgraded nodes, so that services are not affected by the upgrade.

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 Upgrades 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 without user intervention.
- Joining Inaccessible Nodes After Forcing an Upgrade
 You can add inaccessible nodes to the cluster after a forced 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.

Related Topics

<u>Understanding Out-of-Place and Rolling Upgrades</u>
 You can use rolling upgrade or out-of-place upgrades to upgrade your Oracle Grid Infrastructure software.

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.

If you previously attempted to upgrade to a higher release of Oracle Grid Infrastructure, and that Grid home still exists, then instead of running the installer, run the config.bat script in the $Grid_home \cs.$ config directory of the higher version Grid home. Select the Upgrade option to upgrade the existing Grid home to the later release.

1. As the Grid installation user, download the Oracle Grid Infrastructure image files and extract the files to the Grid home.



For example:

md C:\app\21.0.0\grid cd C:\app\21.0.0\grid unzip -q download location\grid home.zip

where download_location\grid_home.zip is the path of the downloaded Oracle Grid Infrastructure image file.

(i) 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.
- If there are non-clustered, or single-instance, Oracle databases that use Oracle ASM running on any of the nodes in the cluster, they must be shut down before you start the upgrade.

Listeners associated with those databases do not have to be shut down.



(i) Note

Oracle recommends that you leave Oracle Real Application Clusters (Oracle RAC) instances running during the Oracle Clusterware upgrade. During the upgrade process, the database instances on the node being upgraded are stopped and started automatically during the upgrade process.

Start the Oracle Grid Infrastructure wizard in the new Grid home by running the following command:

Grid_home\setup.exe

4. Select the option **Upgrade Oracle Grid Infrastructure**.

This option upgrades Oracle Grid Infrastructure (Oracle Clusterware and Oracle ASM).

- On the node selection page, select all nodes. 5.
- **6.** Select installation options as prompted.
- 7. Confirm your selections, and then the upgrade scripts are run automatically.
- 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 or other files that reside in the Oracle Clusterware and Oracle ASM homes.



Note

- Starting from Oracle Grid Infrastructure 12c release 2 (12.2), the OCR and OCR backup must be located in an Oracle ASM disk group. During the upgrade, regardless of where the OCR backup location was in the previous release, the OCR backup location is changed to an Oracle ASM disk group.
- The backups in the old Oracle Clusterware home can be deleted after the
 upgrade because upgrades of Oracle Clusterware are out-of-place upgrades,
 and after the upgrade to the current release of Oracle Grid Infrastructure, the
 OCR and OCR backup must be located in an Oracle ASM disk group.
- 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 without user intervention.

Because the upgrade did not complete successfully on the unreachable nodes, the upgrade is incomplete. Oracle Clusterware remains in the earlier release.

1. Confirm that the upgrade is incomplete by entering the following command:

crsctl query crs activeversion

2. To resolve the incomplete upgrade, run the setup.exe -upgrade command with the -force option on any of the nodes where the setup.exe script has already completed as follows:

 $Grid_home \ crs \ config \ exe - upgrade - force$

For example, as the Oracle Installation User for Oracle Grid Infrastructure, run the following command:

C:\> C:\app\21.0.0\grid\crs\config\setup.exe -upgrade -force

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.

This command forces the upgrade to complete.

3. Verify that the upgrade has completed by using the command crsctl query crs activeversion.

The active release should be the upgrade release.

Joining Inaccessible Nodes After Forcing an Upgrade

You can add inaccessible nodes to the cluster after a forced cluster upgrade.

Starting with Oracle Grid Infrastructure 12c release 1 (12.1), after you complete a force cluster upgrade command, you can join inaccessible nodes to the cluster as an alternative to deleting the nodes, which was required in earlier releases.

To use this option, you must already have Oracle Grid Infrastructure 21c software installed on the nodes.

- 1. Log in as an Administrator user on the nodes that you want to join to the cluster.
- 2. Change directory to the Oracle Grid Infrastructure 21c *Grid_home* directory.

For example:

 $C:\> cd C:\app\21.0.0\grid$

3. Run the following command, where upgraded_node is one of the cluster nodes that is upgraded successfully:

C:\..\grid\> rooters.bat -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.

• **Installation**: If gridconfig.bat fails to complete on the first node, run the following command on another node using the -force option:

Grid_home\crs\config\gridconfig.bat -force -first

• **Upgrade**: If gridconfig.bat fails to complete on the first node, run the following command on another node using the -force option:

Grid_home\crs\config\gridconfig.bat -upgrade -force -first



Upgrading Grid Infrastructure Management Repository from an **Earlier Release**

Complete this procedure to upgrade Grid Infrastructure Management Repository (GIMR) from an earlier release.

You must complete the upgrade of Oracle Grid Infrastructure to the same version before upgrading the GIMR.

As the grid user, create GIMR home directory on all cluster nodes:

 $C:\$ md $C:\$ pp\21.0.0\gimr

Download the Oracle Database image files and extract the files to the GIMR home.

For example:

 $C:\> cd C:\app\21.0.0\gimr$ unzip -q download_location\db_home.zip

download location\db home.zip is the path of the downloaded Oracle Database image file.



- You must extract the image software into the directory where you want your GIMR home to be located.
- Download and copy the Oracle Database image files to the local node only. During upgrade, the software is copied and installed on all other nodes in the cluster.
- From the GIMR home directory, start the Oracle Database software installation:

 $C:\> cd C:\app\21.0.0\gimr$ C:\> setup.exe



(i) Note

Run the setup.exe command from the GIMR home directory only. Do not run setup.exe command that resides at %GIMR_HOME%\oui\bin/, or any other location, to install Oracle Real Application Clusters.

- In the Select Configuration Option screen, select the **Setup Software Only** option. Click Next.
- In the Select Database Installation Option screen, select the Oracle Real Application Clusters Database Installation option. Click Next.
- In the Node Selection screen, select all the nodes where you want to install Oracle RAC or Oracle RAC One Node software. Click Next.
- In the Select Database Edition screen, select the **Enterprise Edition** option. Click **Next**.



- In the Specify Install Location screen, provide the location of Oracle base for Oracle RAC software installation.
 - The Oracle base directory must be different from the GIMR home directory. Click Next.
- 9. The Perform Prerequisite Checks screen displays the results of the prerequisites checks. If any of the checks have a status of Failed and are not Fixable, then you must manually correct these issues. After you have fixed the issue, you can click the Check Again button to have the installer recheck the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click Next.
- 10. Review the contents of the Summary screen and then click Install.
 - The installer displays a progress indicator enabling you to monitor the installation process.
- **11.** Complete GIMR upgrade process using the following command:

C:\> C:\app\21.0.0\gimr\bin\mgmtua upgrade gimrHome

Applying Patches to Oracle Grid Infrastructure

After you have upgraded to Oracle Grid Infrastructure 21c, you can install individual software patches by downloading them from My Oracle Support.

- About Individual Oracle Grid Infrastructure Patches
 Download Oracle ASM individual (one-off) patch and apply it to Oracle Grid Infrastructure using the OPatchAuto utility.
- About Oracle Grid Infrastructure Software Patch Levels
 Review this topic to understand how to apply patches for Oracle ASM and Oracle Clusterware.
- Patching Oracle Grid Infrastructure
 Rolling Patch mode is similar to the existing Rolling Upgrade mode in terms of the Oracle
 ASM operations allowed in this quiesce state.

(i) 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 OPatchAuto utility.

Individual patches are called **one-off** patches. An Oracle ASM one-off patch is available for a specific release of Oracle ASM. If a patch you want is available, then you can download the patch and apply it to Oracle ASM using the OPatchAuto Utility. The OPatch inventory keeps track of the patches you have installed for your release of Oracle ASM. If there is a conflict between the patches you have installed and patches you want to apply, then the OPatchAuto Utility advises you of these conflicts.

About Oracle Grid Infrastructure Software Patch Levels

Review this topic to understand how to apply patches for Oracle ASM and Oracle Clusterware.



The software patch level for Oracle Grid Infrastructure represents the set of all one-off patches and Release Updates (RUs) applied to the Oracle Grid Infrastructure software release, including Oracle ASM. The release is the release number, in the format of major and RU release number. For example, with the release number 21.3.0.0.0, the major release is 21 and the RU release number is 3. The RU release number changes for RU patches, but it does not change for one-off patches.

As with standard upgrades to Oracle Grid Infrastructure, at any given point in time for normal operation of the cluster, all the nodes in the cluster must have the same software release and patch level. Because one-off patches can be applied as rolling upgrades, all possible patch levels on a particular software release are compatible with each other.

Patching Oracle Grid Infrastructure

Rolling Patch mode is similar to the existing Rolling Upgrade mode in terms of the Oracle ASM operations allowed in this quiesce state.

- 1. Download the patch you want to apply from My Oracle Support:
 - a. Go to https://support.oracle.com
 - **b.** Select the Patches and Updates tab to locate the patch.

To locate patch bundles, you can perform a **Product or Family (Advanced)** search for your platform and software release.

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 patches and critical patch updates (CPUs).

- **c.** Place the patch in an accessible directory, such as C:\downloads.
- 2. Change directory to the OPatch directory in the Grid home.

For example:

 $C:\> cd C:\app\21.0.0\grid\opatch$

- 3. Review the patch documentation for the patch you want to apply, and complete all required steps before starting the patch upgrade.
- **4.** Follow the instructions in the patch documentation to apply the patch.

Applying Patches to Grid Infrastructure Management Repository

You can patch Grid Infrastructure Management Repository by installing in-place or out-of-place patches using the OPatch utility.

- In-Place Grid Infrastructure Management Repository Patching
 Use this procedure on each cluster node to perform in-place patching of Grid Infrastructure
 Management Repository (GIMR).
- Out-of-Place Grid Infrastructure Management Repository Patching
 Use this procedure on each cluster node to perform out-of-place patching of Grid
 Infrastructure Management Repository (GIMR).
- In-Place Grid Infrastructure Management Repository Patch Rollback
 Use this procedure on each cluster node for in-place Grid Infrastructure Management
 Repository (GIMR) patch rollback, in case of a failure or an error.



Out-of-Place Grid Infrastructure Management Repository Patch Rollback Use this procedure on each cluster node for out-of-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

In-Place Grid Infrastructure Management Repository Patching

Use this procedure on each cluster node to perform in-place patching of Grid Infrastructure Management Repository (GIMR).

Download patches that 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.

- Review the patch documentation for the patch you want to apply, and complete all required steps.
- As the grid user, run the following command:

C:\> %ORACLE HOME%\bin\mgmtprepatchbootstrap.pl apply inplace -patchHome BOOTSTRAP LOCATION -sourceHome source GIMR home

Where, BOOTSTRAP LOCATION is the directory where the patch has been downloaded and source GIMR home is the Oracle home location.



Note

Do not use slash (\) at the end of the path.

Change directory to the \OPatch directory in the target GIMR home. For example:

C:\> cd GIMR home\OPatch

Run the opatch apply command. For example:

C:\> .\opatch apply patch directory location\patch ID

Run the following command to apply the SQL patch:

C:\> %ORACLE_HOME%\bin\mgmtca applyDatapatch -postPatchApply -id patchIDList

Where, patchIDList is the list of GIMR patch IDs that you want to apply. The -id flag is optional.

Out-of-Place Grid Infrastructure Management Repository Patching

Use this procedure on each cluster node to perform out-of-place patching of Grid Infrastructure Management Repository (GIMR).



Download patches that 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.

- Review the patch documentation for the patch you want to apply, and complete all required steps.
- Change directory to the \OPatch directory in the target GIMR home. For example:

C:\> cd GIMR home\OPatch

Run the opatch apply command on each cluster node in the target GIMR home. For example:

C:\> .\opatch apply patch_directory_location\patch_ID -sourceHome source_GIMR_home

Where, source GIMR home is the Oracle home location.



Note

Do not use slash (\) at the end of the source GIMR home path.

Run the following command on the last node in the target GIMR home:

C:\> mgmtca applyDatapatch -postPatchApply -id patchIDList -sourceHome source_GIMR_home

Where, patchIDList is the list of GIMR patch IDs that you want to apply and source_GIMR_home is the Oracle home location.

The -id flag is optional.

In-Place Grid Infrastructure Management Repository Patch Rollback

Use this procedure on each cluster node for in-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

- Review the patch documentation for the patch you want to roll back, and complete all required steps before starting the rollback.
- As the grid user, run the following command:

C:\> mgmtca applyDatapatch -prePatchRollback -id patchIDList

Where, patchIDList is the list of GIMR patch IDs that you want to roll back.

Change directory to the \OPatch directory in the GIMR home. For example:

C:\> cd GIMR_home\OPatch



4. Run the *opatch -rollback* command. For example:

C:\> .\opatch -rollback directory_location\patch_ID

5. Run the following command to roll back the SQL patch:

C:\> mgmtca applyDatapatch -postPatchRollback -id patchIDList -sourceHome prepatch_home_location

Where, patchIDList is the list of GIMR patch IDs that you want to roll back.

The -sourceHome flag is optional.

Out-of-Place Grid Infrastructure Management Repository Patch Rollback

Use this procedure on each cluster node for out-of-place Grid Infrastructure Management Repository (GIMR) patch rollback, in case of a failure or an error.

- 1. Review the patch documentation for the patch you want to roll back, and complete all required steps before starting the rollback.
- 2. Change directory to the \OPatch directory in the cloned GIMR home. For example:

C:\> cd GIMR_home\OPatch

3. Run the *opatch -rollback* command. For example:

C:\> .\opatch -rollback directory location\patch ID

4. Run the following command from *patchHome*\bin:

C:\> mgmtca applyDatapatch -postPatchRollback -sourceHome source_GIMR_home

Where, *patchHome* is the directory where the patch has been downloaded and *source_GIMR_home* is the pre-patched Oracle home location.

Updating Oracle Enterprise Manager Cloud Control Target Parameters

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

Because Oracle Grid Infrastructure 21c is an out-of-place upgrade of the Oracle Clusterware home in a new location (the Oracle Grid Infrastructure for a cluster home, or Grid home), the path for the CRS_HOME parameter in some parameter files must be changed. If you do not change the parameter, then you encounter errors such as "cluster target broken" on Oracle Enterprise Manager Cloud Control.

To resolve the issue, update the Enterprise Manager Cloud Control target, and then update the Enterprise Manager Agent Base Directory on each cluster member node running an agent.

Updating the Enterprise Manager Cloud Control Target After Upgrades
 After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.



- Updating the Enterprise Manager Agent Base Directory After Upgrades
 After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Agent Base
 Directory on each cluster member node running an agent.
- Registering Resources with Oracle Enterprise Manager After Upgrades
 After upgrading Oracle Grid Infrastructure, add the new resource targets to Oracle Enterprise Manager Cloud Control.

Updating the Enterprise Manager Cloud Control Target After Upgrades

After upgrading Oracle Grid Infrastructure, update the Enterprise Manager Target with the new Grid home path.

- 1. Log in to Enterprise Manager Cloud Control.
- 2. Navigate to the **Targets** menu, and then to the **Cluster** page.
- 3. Click a cluster target that was upgraded.
- 4. Click Cluster, then Target Setup, and then Monitoring Configuration from the menu.
- 5. Update the value for **Oracle Home** with the new Grid home path.
- 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 C:\app\emagent, then Oracle creates the Management Agent home as C:\app\emagent\core\13.1.1.0.

- 1. Navigate to the bin directory in the Management Agent home.
- 2. In the C:\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.
- 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.



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.
- Click Next.
- 7. In the Target Discovery: Results page, select the discovered Oracle ASM Listener target, and click Configure.
- In the Configure Listener dialog box, specify the listener properties and click **OK**.
- 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.

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).

Review your CHM repository needs, and determine if you need to increase the repository size to maintain a larger CHM repository.



(i) 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.

To enlarge the CHM repository, use the following command syntax, where **RETENTION** TIME is the size of CHM repository in number of seconds:

oclumon manage -repos changeretentiontime RETENTION TIME



For example, to set the repository size to four hours:

oclumon manage -repos changeretentiontime 14400

The value for *RETENTION_TIME* must be more than 3600 (one hour) and less than 259200 (three days). If you enlarge the CHM repository size, then you must ensure that there is local space available for the repository size you select on each node of the cluster. If you do not have sufficient space available, then you can move the repository to shared storage.

Downgrading Oracle Clusterware to an Earlier Release

After a successful or a failed upgrade, you can downgrade Oracle Clusterware to the previous release.

- About Oracle Grid Infrastructure Downgrades
 Understand the process and impacts of Oracle Grid Infrastructure downgrades.
- Options for Oracle Grid Infrastructure Downgrades
 You can downgrade Oracle Grid Infrastructure to earlier releases after a successful or a failed upgrade.
- Restrictions for Oracle Grid Infrastructure Downgrades
 Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.
- <u>Downgrading Oracle Standalone Cluster</u>
 Use this procedure to downgrade Oracle Standalone Cluster after a successful upgrade.
- Downgrading Oracle Grid Infrastructure when Upgrade Fails
 If upgrade of Oracle Grid Infrastructure fails before CVU post upgrade checks succeed, then you can run setup.exe and downgrade Oracle Grid Infrastructure to an earlier release.

About Oracle Grid Infrastructure Downgrades

Understand the process and impacts of Oracle Grid Infrastructure downgrades.

Downgrading Oracle Clusterware restores the Oracle Clusterware configuration to the state it was in before the Oracle Grid Infrastructure upgrade. Any configuration changes you performed during or after the Oracle Grid Infrastructure upgrade are removed and cannot be recovered.

To restore Oracle Clusterware to the previous release, use the downgrade procedure for the release to which you want to downgrade.

(i) Note

- You can downgrade the cluster nodes in any sequence. You can downgrade all cluster nodes except one, in parallel. You must downgrade the last node after you downgrade all other nodes.
- When downgrading after a failed upgrade, if the rooters.sh or rooters.bat file does not
 exist on a node, then instead of executing the script, use the command perl
 rooters.pl. Use the Perl interpreter located in the Oracle Home directory.



Options for Oracle Grid Infrastructure Downgrades

You can downgrade Oracle Grid Infrastructure to earlier releases after a successful or a failed upgrade.

Downgrade options include the following earlier releases:

- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 19c.
- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 18c.
- Oracle Grid Infrastructure downgrade to Oracle Grid Infrastructure 12c Release 2 (12.2).

(i) Note

When you downgrade Oracle Grid Infrastructure to an earlier release, for example from Oracle Grid Infrastructure 21c to Oracle Grid Infrastructure 19c, the later release RAC databases already registered with Oracle Grid Infrastructure will not start after the downgrade.

Restrictions for Oracle Grid Infrastructure Downgrades

Review the following information for restrictions and changes for downgrading Oracle Grid Infrastructure installations.

 You can only downgrade to the Oracle Grid Infrastructure release you upgraded from. For example, if you upgraded from Oracle Grid Infrastructure 12c Release 2 (12.2) to Oracle Grid Infrastructure 21c, you can only downgrade to Oracle Grid Infrastructure 12c Release 2 (12.2).

Downgrading Oracle Standalone Cluster

Use this procedure to downgrade Oracle Standalone Cluster after a successful upgrade.

1. As the grid user, delete the Oracle Grid Infrastructure 21c Management Database:

%ORACLE_HOME%\bin\dbca -silent -deleteDatabase -sourceDB -MGMTDB

2. As the grid user, use the rooters.bat -downgrade command from 21c Grid home to downgrade Oracle Grid Infrastructure on all nodes, in any sequence. For example:

%ORACLE_HOME%\crs\install\rootcrs.bat -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 the grid user, downgrade the last node after you downgrade all other nodes:

%ORACLE HOME%\crs\install\rootcrs.bat -downgrade

4. Remove Oracle Grid Infrastructure 21c Grid home as the active Oracle Clusterware home:



- a. On any of the cluster member nodes where the rootupgrade.bat command has run successfully, log in as the Oracle Grid Infrastructure installation owner.
- **b.** Use the following command to start the installer, where C:\app\21.0.0\grid is the location of the new (upgraded) Grid home:

```
cd %ORACLE_HOME%\oui\bin
setup.exe -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=false
ORACLE_HOME=C:\app\21.0.0\grid
"CLUSTER NODES=node1,node2,node3"
```

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

- 5. Set previous release Grid home as the active Oracle Clusterware home:
 - a. On any of the cluster member nodes where the rootupgrade.bat command 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
setup.exe -nowait -waitforcompletion -ignoreSysPrereqs
-updateNodeList -silent CRS=true
ORACLE_HOME=pre-upgrade-grid-home
"CLUSTER_NODES=node1,node2,node3"
```

pre-upgrade-grid-home is the previous release Grid home.

6. As the grid user, start the Oracle Clusterware stack on all nodes from previous release Grid home:

crsctl start crs

7. As the grid user, remove the MGMTDB resource from previous Oracle Grid Infrastructure node:

```
%ORACLE_HOME%\bin\srvctl remove mgmtdb -f
```

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

```
%ORACLE_HOME%\bin\dbca -silent -createDatabase -createAsContainerDatabase true -templateName MGMTSeed_Database.dbc -sid -MGMTDB -gdbName _mgmtdb -storageType ASM -diskGroupName ASM_DG_NAME -datafileJarLocation template-data-file-location -characterset AL32UTF8 -autoGeneratePasswords -skipUserTemplateCheck
```

template-data-file-location is the complete path of the template data file.

9. Configure the Management Database by running the Configuration Assistant from the %ORACLE_HOME%\bin\mgmtca.bat —local location.



Downgrading Oracle Grid Infrastructure when Upgrade Fails

If upgrade of Oracle Grid Infrastructure fails before CVU post upgrade checks succeed, then you can run setup.exe and downgrade Oracle Grid Infrastructure to an 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.bat in silent mode, to downgrade Oracle Clusterware:

C:\> C:\app\21.0.0\grid\grid\Setup.bat -silent -downgrade -nodes nodes_to_be_downgraded -oldHome previous_release_grid_home_to_downgrade_to

Note

- You can downgrade the cluster nodes in any sequence.
- Ensure that the user who performs the downgrade procedure is the same user who performed the upgrade procedure, and belongs to the same user groups, ORA_ASMDBA and ORA_ASMADMIN.

Completing Failed or Interrupted Installations and Upgrades

If Oracle Universal Installer (OUI) exits on the node from which you started the installation or upgrade (the *first node*), or the node reboots before you confirm that the <code>gridconfig.bat</code> script was run on all cluster nodes, then the upgrade or installation 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 Nodes
 When the first node cannot be upgraded, use these steps to continue the upgrade process.
- Continuing Incomplete Upgrades on Remote Nodes
 For nodes other than the first node (the node on which you started the upgrade), use these steps to continue the upgrade process.
- Continuing Incomplete Installations on First Nodes
 To continue an incomplete installation, the first node must finish before the rest of the clustered nodes.
- Continuing Incomplete Installation on Remote Nodes
 For nodes other than the first node (the node on which you started the installation), use these steps to continue the installation process.



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 you click **Ignore** to continue the installation, then the installation or upgrade completes on only some nodes in the cluster. Follow this procedure to add the failed nodes to the cluster.

1. Remove the Oracle Grid Infrastructure software from the failed nodes:

C:\>cd Grid_home\deinstall\deinstall -local

2. As the installation user for Oracle Grid Infrastructure, from a node where Oracle Clusterware is installed, delete the failed nodes from the cluster using the crsctl delete node command:

C:\>cd Grid_home\bin\crsctl delete node -n node_name

where node_name is the node to be deleted.

Run the Oracle Grid Infrastructure installation wizard.

C:\>cd *Grid_home*\setup.exe

After the installation wizard starts, choose the option **Add more nodes to the cluster**, then follow the steps in the wizard to add the nodes.

Alternatively, you can also add the nodes by running the addnode script and specifying the nodes to add to the cluster:

C:\>cd Grid home\addnode\addnode.bat

The failed nodes are added to the cluster.

Continuing Incomplete Upgrade of First Nodes

When the first node cannot be upgraded, use these steps to continue the upgrade process.

- 1. If the OUI failure indicated a need to reboot by raising error message CLSRSC-400, then reboot the first node (the node on which you started the upgrade). Otherwise, manually fix or clear the error condition, as reported in the error output.
- 2. Log in as an Administrator user on the first node.
- Change directory to the new Grid home on the first node, and run the setup.exe —upgrade command on that node again. For example:

C:\> cd app\21.0.0\grid\crs\config\ C:\..\grid> setup.exe -upgrade

4. Complete the upgrade of all other nodes in the cluster.

C:\app\21.0.0\grid\crs\config> setup.exe -upgrade



- Configure a response file, and provide passwords required for the upgrade.
- To complete the upgrade, log in to the first node as the Oracle Installation user for Oracle Grid Infrastructure and run the script setup.exe, located in the Grid home, specifying the response file that you created.

For example, if the response file is named gridinstall.rsp:

C:\> cd *Grid home*\crs\config

C:\Grid home> setup.exe -executeConfigTools -responseFile Grid home\install\response\gridinstall.rsp



(i) Note

When you reupgrade Oracle Grid Infrastructure, you must use the -all flag with the executeConfigTools command to run all the configuration tools.

Related Topics

Postinstallation Configuration Using the ConfigToolAllCommands Script

Continuing Incomplete Upgrades on Remote Nodes

For nodes other than the first node (the node on which you started the upgrade), use these steps to continue the upgrade process.

- If the OUI failure indicated a need to reboot, by raising error message CLSRSC-400, then reboot the node with the error condition. Otherwise, manually fix or clear the error condition that was reported in the error output.
- On the first node, within OUI, click **Retry**.

This instructs OUI to retry the upgrade on the affected node.

3. Continue the upgrade from the OUI instance on the first node.

Continuing Incomplete Installations on First Nodes

To continue an incomplete installation, the first node must finish before the rest of the clustered nodes.

- If the OUI failure indicated a need to reboot, by raising error message CLSRSC-400, then reboot the first node, which is the node where the installation was started. Otherwise, manually fix or clear the error condition that was reported in the error output.
- 2. If necessary, log in as the Oracle Installation user for Oracle Grid Infrastructure. Change directory to the Grid home on the first node and run the setup.exe command on that node again.

For example:

C:\> cd app\21.0.0\grid\crs\config\ C:\..\config> setup.exe

- Complete the installation on all other nodes.
- Configure a response file, and provide passwords for the installation.



5. To complete the installation, log in as the Oracle Installation user for Oracle Grid Infrastructure, and run the following command:

setup.exe --executeConfigTools -responseFile CRS_home/install/response/gridinstall.rsp

For example, if the response file is named gridinstall.rsp:

 $\label{constraint} $C:\> cd\ app\21.0.0\grid\ setup.exe --executeConfigTools -gridinstall.rsp\ CRS_home/install/response/gridinstall.rsp$

Related Topics

<u>Postinstallation Configuration Using the ConfigToolAllCommands Script</u>
 You can create and run a response file configuration after installing Oracle software.

Continuing Incomplete Installation on Remote Nodes

For nodes other than the first node (the node on which you started the installation), use these steps to continue the installation process.

- 1. If the OUI failure indicated a need to reboot, by raising error message CLSRSC-400, then reboot the node with the error condition. Otherwise, manually fix or clear the error condition that was reported in the error output.
- On the first node, within OUI, click Retry.
- 3. Continue the installation from the OUI instance on the first node.

Modifying or Deinstalling Oracle Grid Infrastructure

You must follow a specific procedure when modifying or removing Oracle Clusterware and Oracle Automatic Storage Management (Oracle ASM) software.

- <u>Deciding When to Deinstall Oracle Clusterware</u>
 There are certain situations in which you might be required to remove Oracle software.
- Migrating Standalone Grid Infrastructure Servers to a Cluster

 If you have an Oracle Database installation using Oracle Restart (an Oracle Grid
 Infrastructure installation for a standalone server), then you can reconfigure that server as a cluster member node. Complete the following tasks.
- Changing the Oracle Grid Infrastructure Home Path
 After installing Oracle Grid Infrastructure for a cluster, you might need to change the location of the Grid home.
- Unconfiguring Oracle Clusterware Without Removing the Software
 You can unconfigure Oracle Clusterware on those nodes where you encounter an
 installation error by running the rootcrs.bat -deconfig -force command.
- Removing Oracle Clusterware and Oracle ASM Software
 The deinstall command removes Oracle Clusterware and Oracle ASM from your server.

(i) See Also

Product-specific documentation for requirements and restrictions to remove an individual product

Deciding When to Deinstall Oracle Clusterware

There are certain situations in which you might be required to remove Oracle software.

Remove installed components in the following situations:

- You have successfully installed Oracle Clusterware, and you want to remove the Oracle Clusterware installation, either in an educational environment, or a test environment.
- You have encountered errors during or after installing or upgrading Oracle Clusterware, and you want to reattempt an installation.
- Your installation or upgrade stopped because of a hardware or operating system failure.
- You are advised by Oracle Support to reinstall Oracle Clusterware.



Migrating Standalone Grid Infrastructure Servers to a Cluster

If you have an Oracle Database installation using Oracle Restart (an Oracle Grid Infrastructure installation for a standalone server), then you can reconfigure that server as a cluster member node. Complete the following tasks.

 Inspect the Oracle Restart configuration with the Server Control (SRVCTL) utility using the following syntax, where db_unique_name is the unique name for the database, and Isnrname is the name of the listener for the database:

```
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; you will need this information in a later step.

- 2. Stop all of the databases, services, and listeners that you discovered in step 1.
- 3. Log in as an Administrator user and navigate to the directory *Grid_home*\crs\install, where *Grid_home* is the location of your Oracle Grid Infrastructure home (Grid home) directory, for example:

C:\> cd app\21.0.0\grid\crs\install

4. Unconfigure the Oracle Grid Infrastructure installation for a standalone server (Oracle Restart) using the following command:

C:\..\install> roothas.bat -deconfig -force

5. Prepare the server for Oracle Clusterware configuration, as described in Chapter 2 through Chapter 7 of this guide. In addition, choose to install Oracle Grid Infrastructure for a cluster in the same location as Oracle Restart, or in a different location:

Option	Description	
Installing in the Same Location as Oracle Restart	Proceed to step 7.	
Installing in a Different Location than Oracle Restart	Set up Oracle Grid Infrastructure software in the new Grid home software location, then proceed to step 7.	

6. 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
export oracle_install_asm_ConfigureGIMRDataDG=true or false
export oracle_install_asm_GIMRDataDGName=disk_group_name
```

As the Oracle Installation user for Oracle Grid Infrastructure, run the installer.

You can complete the installation interactively, or if you want to perform a silent installation, then save and stage the response file. After saving the response file, run the following



command. For the -responseFile parameter, specify the full path name where the response file was saved, for example:

C:\> *Grid_home*\setup.exe -silent -responseFile C:\Users\dba1\scripts\GI.rsp

- 8. Mount the Oracle ASM disk group used by Oracle Restart.
- 9. 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 -pwfile -oraclehome %ORACLE_HOME% -node *nodename*

- a. Verify that the %ORACLE_HOME% environment variable is set to the location of the database home directory.
- **b.** To add the database name mydb, enter the following command:

srvctl add database -db mydb -spfile -pwfile -oraclehome %ORACLE_HOME% -node node1

c. Add each service to the database, using the command srvctl add service. For example:

srvctl add service -db mydb -service myservice

10. Add nodes to your cluster, as required, using the Oracle Grid Infrastructure installer.

(i) See Also

- Recording Response Files for more information about saving the response file
- Oracle Clusterware Administration and Deployment Guide for information about adding nodes to your cluster

Changing the Oracle Grid Infrastructure Home Path

After installing Oracle Grid Infrastructure for a cluster, you might need to change the location of the Grid home.

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.

Before changing the Grid home, you must shut down all executables that run in the Grid home directory that you are modifying. In addition, shut down all applications that use Oracle shared libraries.

1. Login as an Administrator user or the Oracle Installation user for Oracle Grid Infrastructure (for example, grid).



2. Change directory to *Grid home*\bin and enter the command crsctl stop crs. For example:

C:\> cd app\21.0.0\grid\bin C:\..\BIN> crsctl stop crs

3. Detach the existing Grid home.

Run a command similar to the following command, where $C:\app\21.0.0\grid$ is the existing Grid home location:

C:\> cd app\21.0.0\grid\oui\bin
C:\..\bin> setup.exe -silent -detachHome ORACLE_HOME= 'C:\app\21.0.0\grid' -local

Move the installed files for Oracle Grid Infrastructure from the old Grid home to the new Grid home.

For example, if the old Grid home is C:\app\21.0.0\grid and the new Grid home is D:\app\21c\grid, use the following command:

C:\> xcopy C:\app\21.0.0\grid D:\app\21c\grid /E /I /H /K

Install Oracle Grid Infrastructure in the new home.

For example:

C:\> perl clone.pl ORACLE_BASE=C:\app\grid ORACLE_HOME=C:\app\21.0.0\grid ORACLE_HOME_NAME=OraHome1Grid "LOCAL_NODE=node1" "CLUSTER_NODES={node1,node2}" CRS=TRUE

When you navigate to the *Grid_home*\clone\bin directory and run the clone.pl script, provide values for the input parameters that provide the path information for the new Grid home.

6. Start Oracle Clusterware in the new home location.

 $\label{eq:cdapp} $D:\> cd app\21c\grid\crs\install $$D:\..install\> rooters.bat -dstcrshome $D:\app\21c\grid -move $$$



🛕 Warning

While cloning, ensure that you do not change the Oracle home base, otherwise the move operation will fail.

7. Repeat steps 1 through 6 on each cluster member node.

Related Topics

Oracle Clusterware Administration and Deployment Guide



Unconfiguring Oracle Clusterware Without Removing the Software

You can unconfigure Oracle Clusterware on those nodes where you encounter an installation error by running the rooters.bat -deconfig -force command.

After you unconfigure, correct the cause of the error, and then run rooters.bat again to reconfigure Oracle Clusterware.

Running the rooters.bat command with the options -deconfig -force enables you to unconfigure Oracle Clusterware on one or more nodes without removing the installed software. This feature is useful if you encounter an error on one or more cluster nodes during installation, such as incorrectly configured shared storage.

Before unconfiguring Oracle Clusterware you must stop any databases, services, and listeners that may be installed and running.

Commands used in this section will remove the Oracle Grid infrastructure installation for the entire cluster. To remove the installation from an individual node, refer to Oracle Clusterware Administration and Deployment Guide.

- Log in using a member of the Administrators group on a node where you encountered an error during installation.
- Stop any databases, services, and listeners currently running from the Grid home.
- Change directory to *Grid_home*\crs\install.

For example:

C:\> cd C:\app\21.0.0\grid\crs\install

Run rooters.bat with the -deconfig -force options.

For example:

C:\..\install> rooters.bat -deconfig -force

(i) Note

The -force option must be specified when running the rooters.bat script if there exist running resources that depend on the resources started from the Oracle Clusterware home you are deleting, such as databases, services, or listeners. You must also use the -force option if you are removing a partial, or failed installation.

- Repeat Step 1 through Step 4 on other nodes as required.
- If you are unconfiguring Oracle Clusterware on all nodes in the cluster, then on the last node, enter the following command:

C:\..\install> rooters.bat -deconfig -force -lastnode



The -lastnode option completes the unconfiguring of the cluster, including the Oracle Cluster Registry (OCR) and voting files.



Run the rooters.bat -deconfig -force -lastnode command on a cluster node.

Removing Oracle Clusterware and Oracle ASM Software

The deinstall command removes Oracle Clusterware and Oracle ASM from your server.

∴ Caution

You must run the deinstall command from the same release to remove Oracle software. Do not run the deinstall command from a later release to remove Oracle software from an earlier release. For example, do not run the deinstall command from the 21c Oracle home to remove Oracle software from an existing 12.2, 18c, or 19c Oracle home.

About Oracle Deinstallation Options

The deinstall.bat command stops Oracle software, and removes Oracle software and configuration files from the operating system.

- Files Deleted by the deinstall Command
 - The deinstall command removes Oracle software and files from your system.
- deinstall Command Reference
 - Use the deinstall command to remove Oracle software. You can run this command from an Oracle home directory after installation.
- Using the deinstall Command to Remove Oracle Clusterware and Oracle ASM
 You can run the deinstall command in multiple ways.

About Oracle Deinstallation Options

The deinstall.bat command stops Oracle software, and removes Oracle software and configuration files from the operating system.

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.



Note

- You must run the deinstall command from the same release to remove Oracle software. Do not run the deinstall command from a later release to remove Oracle software from an earlier release. For example, do not run the deinstall command from the 21c Oracle home to remove Oracle software from an existing 19.0.0 Oracle home.
- After running the deinstall command, you must remove the corresponding Oracle home directory and stop all services.
- Starting with Oracle Database 12c Release 1 (12.1.0.2), the roothas.bat script replaces the roothas.pl script in the Oracle Grid Infrastructure home for Oracle Restart, and the rooters.bat script replaces the rooters.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 the deinstall cannot determine the configuration. You must provide all the configuration details either interactively or in a response file.

Files Deleted by the deinstall Command

The deinstall command removes Oracle software and files from your system.

When you run deinstall, if the central inventory (Inventory) contains no other registered homes besides the home that you are deconfiguring and removing, then the 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

The deinstall command 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.



deinstall Command Reference

Use the deinstall command to remove Oracle software. You can run this 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.

File Path

%ORACLE HOME%\deinstall

Prerequisites

Before you run the deinstall command for Oracle Grid Infrastructure installations, notify your DNS administrator to delete the subdomain entry from the DNS, if Grid Naming Service (GNS) is in use.

Syntax When Using the deinstall.bat Program

deinstall.bat

```
[-silent] [-checkonly] [-paramfile complete path on input parameter properties file]
[-checkonly]
[-local]
[-paramfile complete path of input parameter properties 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]
[-skipLocalHomeDeletion]
[-skipRemoteHomeDeletion]
[-help]
```

Options

Table 12-1 Options for the Deinstallation Tool

Command Option	Description
home complete path of Oracle home	Specify this option to indicate the home path of the Oracle home to check or deinstall. To deinstall Oracle software using the deinstall.bat command located within the Oracle home being removed, provide a response file in a location outside the Oracle home, and do not use the -home option. If you run deinstall.bat from the %ORACLE_HOME% \deinstall path, then the -home option is not required because the command knows from which home it is being run.



Table 12-1 (Cont.) Options for the Deinstallation Tool

Command Option	Description
silent	Specify this option to run the deinstall in noninteractive mode. This option requires one of the following: A working system that it can access to determine the installation and configuration information; the -silent option does not work with failed installations. A response file that contains the configuration
	values for the Oracle home that is being deinstalled or deconfigured.
checkonly	Specify this option to check the status of the Oracle software home configuration. Running the deinstall command with the -checkonly option does not remove the Oracle configuration. This option generates a response file that you can use with the deinstall.bat command.
	When you use the -checkonly option to generate a response file, you are prompted to provide information about your system. You can accept the default value the deinstall command has obtained from your Oracle installation, indicated inside brackets ([]), or you can provide different values. To accept the defaults, press Enter at each prompt.
local	Specify this option on a multinode environment to deconfigure Oracle software in a cluster.
	When you run deinstall.bat with this option, it deconfigures and deinstalls the Oracle software only on the local node (the node on which you run deinstall.bat) for non-shared Oracle home directories. The deinstall command does not deinstall or deconfigure Oracle software on remote nodes.
paramfile complete path of input parameter properties file	(Optional) You can specify this option to run deinstall.bat with a response file in a location other than the default. When you use this option, provide the complete path where the response file is located. If you run the deinstall.bat command from the Oracle home that you plan to deinstall, then you do not need to specify the -paramfile option. The default location of the response file is %ORACLE_HOME%\deinstall\response.
params name1=value[name2=value name3=value]	Use this option with a response file to override one or more values in a response file that you created.
o complete path of directory for saving response files	Use this option to provide a path other than the default location where the response file (deinstall.rsp.tmpl) is saved.
	The default location of the response file is %ORACLE_HOME%\deinstall\response.
tmpdir complete path of temporary directory to use	Specifies a non-default location where the deinstall command writes the temporary files for the deinstallation.



Table 12-1 (Cont.)	Options f	for the I	Deinstallation	Tool
--------------	--------	------------------	-----------	----------------	------

Command Option	Description
logdir complete path of log directory to use	Specifies a non-default location where the deinstall command writes the log files for the deinstallation.
skipLocalHomeDeletion	Specify this option in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a local Grid home without deleting the Grid home.
skipRemoteHomeDeletion	Specify this option in Oracle Grid Infrastructure installations on a multinode environment to deconfigure a remote Grid home without deleting the Grid home.
help	Use the -help option to obtain additional information about the command option flags.

Location of Log Files for the Deinstallation Tool

If you use the deinstall.bat command located in an Oracle home, then the deinstall writes log files in the C:\Program Files\Oracle\Inventory\logs directory.

If you are using the deinstall.bat command to remove the last Oracle home installed on the server, then the log files are written to the current user's home directory. For example, if you are logged in as the domain user RACDBA\dba1, then the log files are stored in the directory C:\Users\dba1.RACDBA\logs.

Using the deinstall Command to Remove Oracle Clusterware and Oracle ASM

You can run the deinstall command in multiple ways.

- Running the deinstall Command From an Oracle Home
 You can run the deinstall command from an Oracle home.
- Running the deinstall command with a Response File
 When running the deinstall command, you can use a response file instead of responding to each prompt individually.
- Generating a Response File For Use With the deinstall Command
 To use a response file with the deinstall command, you must first create the response file.

Running the deinstall Command From an Oracle Home

You can run the deinstall command from an Oracle home.

- 1. The default method for running the deinstall command is from the deinstall directory in the Oracle home as the Oracle Installation user:
 - C:\> %ORACLE HOME%\deinstall\deinstall.bat
- Provide information about your servers as prompted or accept the defaults.

The deinstall command stops Oracle software, and removes Oracle software and configuration files on the operating system.



Example 12-1 Running deinstall.bat From Within the Oracle Home

The most common method of running the deinstall command is to use the version installed in the Oracle home being removed. The deinstall command determines the software configuration for the local Oracle home, and then provides default values at each prompt. You can either accept the default value, or override it with a different value. If the software in the Oracle home is not running (for example, after an unsuccessful installation), then the deinstall command cannot determine the configuration, and you must provide all the configuration details either interactively or in a response file. To use the deinstall command located in the Oracle home directory, issue the following command, where C:\app\\21.0.0\grid is the location of Grid home:

Use the following command while logged in as a member of the Administrators group to remove the Oracle Grid Infrastructure installation from your cluster:

C:\> C:\app\21.0.0\grid\deinstall\deinstall.bat

Provide additional information as prompted.



Note

When using the deinstall command from a location other than within the Oracle home being removed, you must specify the -home option on the command line.

Running the deinstall command with a Response File

When running the deinstall command, you can use a response file instead of responding to each prompt individually.

The deinstall command uses the information you provide, plus information gathered from the software home to create a response file. You can alternatively supply a response file generated previously by the deinstall.bat command using the -checkonly option and -o option. You can also edit a response file template located at Oracle Home\deinstall\response\deinstall.rsp.tmpl to create a response file.

To run the deinstall.bat command located in an Oracle Grid Infrastructure home and use a response file located at D:\Users\oracle\paramfile4.tmpl, enter the following commands while logged in as a member of the Administrators group:

C:\> cd %ORACLE_HOME% C:\..grid> deinstall\deinstall.bat -paramfile D:\Users\oracle\paramfile4.tmpl

Generating a Response File For Use With the deinstall Command

To use a response file with the deinstall command, you must first create the response file.

You can generate the a response file by running the deinstall.bat command with the -checkonly and -o options before you run the command to deinstall the Oracle home, or you can use the response file template and manually edit it to create the response file.

Alternatively, you can use the response file template located at %ORACLE HOME% \deinstall\response\deinstall.rsp.tmpl.

To generate the response file deinstall_OraCrs21c_home1.rsp using the deinstall.bat command located in the Oracle home and the -checkonly option, enter a command similar to the



following, where C:\app\21.0.0\grid is the location of the Grid home and C:\Users\oracle is the directory in which the generated response file is created:

 $C:\> app\21.0.0\grid\deinstall\deinstall.bat\-checkonly\-o\ C:\Users\oracle\$



Installing and Configuring Oracle Grid Infrastructure Using Response Files

You can install and configure Oracle Grid Infrastructure software using response files.

- How Response Files Work
 - Response files can assist you with installing an Oracle product multiple times on multiple computers.
- <u>Deciding to Use Silent Mode or Response File Mode</u>
 There are several reasons for running the installer in silent mode or response file mode.
- Using Response Files
 - Use these general steps for installing and configuring Oracle products using the installer in silent or response file mode.
- Preparing Response Files
 - Review this information to prepare response files for use during silent mode or response file mode installations.
- Running Oracle Universal Installer Using a Response File
 After creating the response file, run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation.
- Running Oracle Net Configuration Assistant Using Response Files
 You can run Oracle Net Configuration Assistant (NETCA) in silent mode to configure and start an Oracle Net listener on the system, configure naming methods, and configure Oracle Net service names.
- <u>Postinstallation Configuration Using Response File Created During Installation</u>
 Use response files to configure Oracle software after installation. You can use the same response file created during installation to also complete postinstallation configuration.
- <u>Postinstallation Configuration Using the ConfigToolAllCommands Script</u>
 You can create and run a response file configuration after installing Oracle software.

How Response Files Work

Response files can assist you with installing an Oracle product multiple times on multiple computers.

When you start the installer, you can use a response file to automate the installation and configuration of Oracle software, either fully or partially. The installer 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, supply the appropriate value for the ORACLE_HOME variable:

ORACLE_HOME=C:\app\oracle\product\21.0.0\dbhome_1

Another way of specifying the response file variable settings is to pass them as command line arguments when you run the installer. For example:

-silent directory_path

In this command, *directory_path* is the path of the database directory on the DVD, or the path of the directory on the hard drive.

Ensure that you enclose the variable and its setting in double-quotes.

Deciding to Use Silent Mode or Response File Mode

There are several reasons for running the installer in silent mode or response file mode.

Table A-1 Reasons for Using Silent Mode or Response File Mode

Mode	Reasons to Use
Silent	 Use silent mode for the following installations: To complete an unattended installation, which you schedule using operating system utilities To complete several similar installations on multiple systems without user interaction Install the software on a system that cannot display the Oracle Universal Installer (OUI) graphical user interface OUI 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 the installer prompts. If you do not specify information required for a particular OUI screen in the response file, then the installer displays that screen. OUI suppresses screens for which you have provided all of the required information.



Using Response Files

Use these general steps for installing and configuring Oracle products using the installer in silent or response file mode.

(i) Note

You must complete all required preinstallation tasks on a system before running the installer in silent or response file mode.

1. Verify the Windows Registry key HKEY_LOCAL_MACHINE\Software\Oracle exists and that the value for inst_loc is the location of the Oracle Inventory directory on the local node.

Note

Changing the value for <code>inst_loc</code> in the Windows registry is not supported after the installation of Oracle software

- 2. Prepare a response file.
- 3. Run the installer in silent or response file mode.
- If you completed a software-only installation, then perform the steps necessary to configure the Oracle product.

Preparing Response Files

Review this information to prepare response files for use during silent mode or response file mode installations.

- About Response File Templates
 - Oracle provides response file templates for each product and installation type and for each configuration tool.
- Editing a Response File Template
 - You can copy and modify a response file template for each product and installation type and for each configuration tool.
- Recording Response Files
 - You can use the installer in interactive mode to record response files, which you can then edit and use to complete silent mode or response file mode installations.

About Response File Templates

Oracle provides response file templates for each product and installation type and for each configuration tool.

For Oracle Database, the response file templates are located in the database\response directory on the installation media and in the *Oracle_home*\inventory\response directory. For Oracle Grid Infrastructure, the response file templates are located in the *Grid_home*\install\response directory after the software is installed.



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 response files are provided with this software:

Table A-2 Response Files for Oracle Database and Oracle Grid Infrastructure

Response File	Used For
db_install.rsp	Silent configuration of Oracle Database software
dbca.rsp	Silent creation and configuration of an Oracle Database using Oracle Database Configuration Assistant (DBCA)
netca.rsp	Silent configuration of Oracle Net using NETCA
grid_install.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 must be given to the Oracle software installation owner only, and access restricted to the response file. Oracle strongly recommends that database administrators or other administrators delete or secure response files when they are not in use.

Editing a Response File Template

You can copy and modify a response file template for each product and installation type and for each configuration tool.

To copy and modify a response file, perform the following steps:

Copy the response file from the response file directory to a directory on your system.

copy Oracle home\install\response\product timestamp.rsp local directory

- Open the response file in a text editor.
- 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.

Secure the response file.

Ensure that only the user that installed the Oracle software can view or modify response files. Consider deleting the modified response file after the installation succeeds.



(i) Note

A fully specified response file for an Oracle Grid Infrastructure installation or an Oracle Database installation can contain the passwords for:

- Oracle Automatic Storage Management (Oracle ASM) administrative accounts
- Database administrative accounts
- A user who is a member of the operating system group ORA DBA (required for automated backups)

Recording Response Files

You can use the installer 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 the installer starts to copy the software to the local disk.

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.

(i) Note

You cannot save passwords while recording the response file.

1. Complete preinstallation tasks 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. Log in as the Oracle Installation User. Ensure that the Oracle Installation User has permissions to create or write to the Grid home path that you specify during installation.
- Start the installer. On each installation screen, specify the required information.
- When the installer displays the Summary screen, perform the following steps:
 - Click Save Response File. In the pop up window, specify a file name and location to save the values for the response file, then click **Save**.
 - **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.





(i) Note

Your response file name must end with the .rsp suffix.

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.

The Oracle Universal Installer executable, setup.exe provides several options. For help information on the full set of these options, run the setup, exe command with the -help option. For example:

For Oracle Database:

```
db home> setup.exe -help
```

For Oracle Grid Infrastructure:

```
Grid_home> setup.exe -help
```

The help information appears in your session window after a short period of time.

To run the installer using a response file, perform the following steps:

- Complete the preinstallation tasks for a normal installation.
- Log in as an Administrator user or the user that installed the software.
- To start the installer in silent or response file mode, enter a command similar to the following:
 - For Oracle Database:

```
C: \ b\_home \ [-silent] \ \ 
   -responseFile response filename
```

For Oracle Grid Infrastructure:

```
C:\> Grid_home\setup.exe [-silent] \
   -responseFile response_filename
```



(i) Note

Do not specify a relative path to the response file. If you specify a relative path, then the installer fails.

In this example:

db_home is the directory on the hard drive where you have copied the Oracle Database installation software.



- Grid_home is the directory on the hard drive where you have copied the Oracle Grid Infrastructure installation software.
- silent runs the installer in silent mode.
- response_filename is the full path and file name of the installation response file that you configured.

Running Oracle Net Configuration Assistant Using Response Files

You can run Oracle Net Configuration Assistant (NETCA) in silent mode to configure and start an Oracle Net listener on the system, configure naming methods, and configure Oracle Net service names.

To run 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 NETCA using a response file:

 Copy the netca.rsp response file template from the response file directory to a directory on your system.

copy \directory_path\assistants\netca\netca.rsp local_directory

In this example, *directory_path* is the path of the directory where you have copied the installation binaries.

If the software is staged on a hard drive, or has already been installed, then you can edit the file in the response directory located on the local disk instead.

- 2. Open the response file in a text editor.
- 3. Follow the instructions in the file to edit it.



NETCA fails if you do not correctly configure the response file.

- 4. Log in as an Administrator user and set the %ORACLE_HOME% environment variable to specify the correct Oracle home directory.
- Enter a command similar to the following to run NETCA in silent mode:

C:\> *Oracle_home*\bin\netca -silent -responsefile *X*:*local_dir*\netca.rsp

In this command:

- The -silent option runs NETCA in silent mode.
- *X:\local_dir* is the full path of the directory where you copied the netca.rsp response file template, where *X* represents the drive on which the file is located, and *local_dir* the path on that drive.



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 Use the response file created during installation to also complete postinstallation configuration.
- Running Postinstallation Configuration Using a Response File Complete this procedure to run configuration assistants configuration with the executeConfigTools command.

Using the Installation Response File for Postinstallation Configuration

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. The owner of the response file must be the installation owner user.

Example A-1 Response File Passwords for Oracle Grid Infrastructure

oracle.install.crs.config.ipmi.bmcPassword=password oracle.install.asm.SYSASMPassword=password oracle.install.asm.monitorPassword=password oracle.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.



(i) Note

IPMI is not available on Microsoft Windows for Oracle Database Release 21c, but will be available in a future release.

If you do not want to enable Oracle Enterprise Manager for management, then leave the emAdminPassword password field blank.

Example A-2 Response File Passwords for Oracle Grid Infrastructure for a Standalone **Server (Oracle Restart)**

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 A-3 Response File Passwords for Oracle Database

This example illustrates the passwords to specify for use with the database configuration assistants.

oracle.install.db.config.starterdb.password.SYS=password oracle.install.db.config.starterdb.password.SYSTEM=password oracle.install.db.config.starterdb.password.DBSNMP=password oracle.install.db.config.starterdb.password.PDBADMIN=password oracle.install.db.config.starterdb.emAdminPassword=password oracle.install.db.config.asm.ASMSNMPPassword=password oracle.install.OracleHomeUserPassword=password

You can also specify oracle.install.db.config.starterdb.password.ALL=password to use the same password for all database users.

Oracle Database configuration assistants require the SYS, SYSTEM, PDBADMIN, and DBSNMP passwords for use with Oracle Database Configuration Assistant (DBCA). Specify the following passwords, depending on your system configuration:

- If the database uses Oracle ASM for storage, then you must specify a password for the ASMSNMPPassword variable. If you are not using Oracle ASM, then leave the value for this password variable blank.
- If you did not specify an Oracle Home user for the Oracle Database installation, then leave the OracleHomeUserPassword field blank.

Running Postinstallation Configuration Using a Response File

Complete this procedure to run configuration assistants configuration with the -executeConfigTools command.

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 Grid
Infrastructure:

oracle.install.asm.SYSASMPassword=*password* oracle.install.config.emAdminPassword=*password*

Change directory to the Oracle home containing the installation software. For example, for Oracle Grid Infrastructure:

cd Grid_home

3. Run the configuration script using the following syntax:

For Oracle Grid Infrastructure:

setup.exe -executeConfigTools -responsefile *Grid_home*\install\response\product_timestamp.rsp



For Oracle Database:

setup.exe -executeConfigTools -responseFile *Oracle_home*\install\response\product_timestamp.rsp

For Oracle Database, you can also edit and use the response file located in the directory *Oracle_home*\inventory\response\:

setup.exe -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:

setup.exe -executeConfigTools -responseFile *Grid_home*\install\response\grid_2016-09-09_01-03-36PM.rsp - silent

For Oracle Database:

setup.exe -executeConfigTools -responseFile *Oracle_home*\inventory\response\db_2016-09-09_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.

Starting with Oracle Database 18c Release, use the <code>executeConfigTools</code> script to complete the postinstall configuration.

- About the Postinstallation Configuration File
 The configuration assistants are started with a script called configToolAllCommands.
- Creating a Password Response File
 Use these steps to create a password response file for use with the configToolAllCommands script
- Running Postinstallation Configuration Using a Script and Response Files
 You can run the postinstallation configuration assistants with the configToolAllCommands script.





"<u>Postinstallation Configuration Using Response File Created During Installation</u>" for an alternate method of postinstallation configuration of Oracle software using the same response file created at the time of installation.

About the Postinstallation Configuration File

The configuration assistants are started with a script called configToolAllCommands.

When you perform an installation using silent mode or response file mode, you provide information about your servers in a response file that you otherwise provide manually using a graphical user interface. However, the response file does not contain passwords for user accounts that configuration assistants require after software installation is complete. To run the configuration assistants after the installation completes in silent mode, you must run the configToolAllCommands script and provide the passwords used by the assistants in a password file.

You can run the <code>configToolAllCommands</code> script in silent mode by using a password response file. The script uses the passwords in the file to run the configuration tools in succession to complete the software configuration. If you keep the password file to use when cloning installations, then Oracle strongly recommends that you store the password file in a secure location.

You can also use the password file to restart a failed installation. If you stop an installation to fix an error, then you can rerun the configuration assistants using configToolAllCommands and a password response file.

The configToolAllCommands password response file has the following 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, to set the password for the SYS user of Oracle ASM:

oracle.crs|S_ASMPASSWORD=PassWord

Creating a Password Response File

Use these steps to create a password response file for use with the configToolAllCommands script..

- 1. Create a response file that has a name of the format *filename*.properties.
- 2. Open the file with a text editor, and cut and paste the sample password file contents, as shown in the example below, modifying as needed.
- If the file is stored on a volume formatted for Windows New Technology File System (NTFS), then modify the security permissions to secure the file.



Example A-4 Sample Password response file for Oracle Grid Infrastructure Installation

Oracle Grid Infrastructure requires passwords for Oracle Automatic Storage Management Configuration Assistant (ASMCA), and for Intelligent Platform Management Interface Configuration Assistant (IPMICA) if you have a baseboard management controller (BMC) card and you want to enable this feature. Provide the following response file:

oracle.crs|S ASMPASSWORD=password oracle.crs|S OMSPASSWORD=password oracle.crs|S ASMMONITORPASSWORD=password oracle.crs|S_BMCPASSWORD=password oracle.crs|S WINSERVICEUSERPASSWORD=password

If you do not have a BMC card, or you do not want to enable Intelligent Platform Management Interface (IPMI), then leave the S BMCPASSWORD input field blank.



(i) Note

IPMI is not available on Microsoft Windows for Oracle Database Release 21c, but will be available in a future release.



(i) Note

If you are upgrading Oracle ASM 11g Release 1 or earlier releases, then you only need to provide the input field for oracle.assistants.asm|S ASMMONITORPASSWORD.

Example A-5 Sample Password response file for Oracle Real Application Clusters

Oracle Database configuration requires the SYS, SYSTEM, and DBSNMP passwords for use with Database Configuration Assistant (DBCA). The S_ASMSNMPPASSWORD response is necessary only if the database is using Oracle ASM for storage. Specify the S PDBADMINPASSWORD password for your multitenant container database (CDB) with one or more pluggable databases (PDBs). Also, if you selected to configure Oracle Enterprise Manager Cloud Control, then you must provide the password for the Oracle software installation owner for S EMADMINPASSWORD, similar to the following example, where the phrase password represents the password string:

oracle.server|S_SYSPASSWORD=password oracle.server|S SYSTEMPASSWORD=password oracle.server|S DBSNMPPASSWORD=password oracle.server|S_PDBADMINPASSWORD=password oracle.server|S_EMADMINPASSWORD=password oracle.server|S ASMSNMPPASSWORD=password oracle.server|S WINSERVICEUSERPASSWORD=password

If you do not want to enable Oracle Enterprise Manager for Oracle ASM, then leave those password fields blank.

Running Postinstallation Configuration Using a Script and Response Files

You can run the postinstallation configuration assistants with the configToolAllCommands script.



- 1. Create a password response file using the format *filename*.properties for the file name, as described in Creating a Password Response File.
- 2. Change directory to *Oracle_home*\cfgtoollogs, and run the configuration script using the following syntax:

configToolAllCommands RESPONSE_FILE=\path\filename.properties

For example:

 $C: \verb|\| Cfgtoollogs| configToolAllCommands RESPONSE_FILE=C: \verb|\| Cfgrsp.properties| | Cfgrs$

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.

Oracle Base Directory Naming Convention

This section describes what the Oracle base is, and how it functions.

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

This topic shows 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.



(i) Note

OFA assists in identification of an ORACLE BASE with its Automatic Diagnostic Repository (ADR) diagnostic data to properly collect incidents.



About Multiple Oracle Homes Support

Oracle Database supports multiple Oracle homes. You can install this release or earlier releases of the software more than once on the same system, in different Oracle home directories.

Careful selection of mount point names can make Oracle software easier to administer. Configuring multiple Oracle homes in compliance with Optimal Flexible Architecture (OFA) rules provides the following advantages:

- You can install this release, or earlier releases of the software, more than once on the same system, in different Oracle home directories. However, you cannot install products from one release of Oracle Database into an Oracle home directory of a different release.
- Multiple databases, of different versions, owned by different users can coexist concurrently.
- To install Oracle Database software in multiple Oracle homes, you must extract the image file in each Oracle home, and then run the setup wizard from the respective Oracle home.
- You must install a new Oracle Database release in a new Oracle home that is separate from earlier releases of Oracle Database.
 - You cannot install multiple releases in one Oracle home. Oracle recommends that you create a separate Oracle Database Oracle home for each release, in accordance with the Optimal Flexible Architecture (OFA) guidelines.
- In production, the Oracle Database server software release is the release number in the format of major and RU release number. For example, with the release number 19.3.0.0.0, the major release is 19 and the RU release number is 3.
- Later Oracle Database releases can access earlier Oracle Database releases. However, this access is only for upgrades. For example, Oracle Database 19c can access an Oracle Database 18c if the 18c database is started up in upgrade mode.
- Oracle Database Client can be installed in the same Oracle Database home if both products are at the same release level. For example, you can install Oracle Database Client 19c into an existing Oracle Database 19c home but you cannot install Oracle Database Client 19c into an existing Oracle Database 18c home. If you apply a patch set or release update before installing the client, then you must apply the patch set or release update again.
- Structured organization of directories and files, and consistent naming for database files simplify database administration.
- Login home directories are not at risk when database administrators add, move, or delete
 Oracle home directories.
- You can test software upgrades in an Oracle home in a separate directory from the Oracle home where your production database is located.
- For information about release support timelines, refer to My Oracle Support Doc ID 742060.1

Related Topics

My Oracle Support Note 742060.1

Oracle Base Directory Naming Convention

This section describes what the Oracle base is, and how it functions.



The Oracle Base directory is the database home directory for Oracle Database installation owners and the log file location for Oracle Grid Infrastructure owners. You should name Oracle base directories using the syntax pmhu, 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 mount point 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 B-1 Examples of OFA-Compliant Oracle Base Directory Names

Example	Description
C:\app\oracle	Oracle base directory for Oracle Database, where the Oracle Database software installation owner name is oracle. The Oracle Database binary home is located underneath the Oracle base path.
C:\app\grid	Oracle base directory for Oracle Grid Infrastructure, where the Oracle Grid Infrastructure software installation owner name is grid.



∧ Caution

The Oracle Grid Infrastructure Oracle base should not contain the Oracle Grid Infrastructure binaries for an Oracle Grid Infrastructure for a cluster installation. Permissions for the file path to the Oracle Grid Infrastructure binary home is changed to LocalSystem or the Oracle Home User, if specified, during installation.

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 $pm\sl v poduct v y pe_n$. The following table describes the variables used in this syntax:

Variable	Description
pm	A mount point name.
S	A standard directory name.
и	The name of the owner of the directory.
V	The version of the software.
type	The type of installation. For example: Database (dbhome), Client (client), or Oracle Grid Infrastructure (grid)



Variable	Description
n	An optional counter, which enables you to install the same product more than once in the same Oracle base directory. For example: Database 1 and Database 2 (dbhome_1, dbhome_2)

For example, the following path is typical for the first installation of Oracle Database on this system:

C:\app\oracle\product\21.0.0\dbhome_1

Optimal Flexible Architecture File Path Examples

This topic shows examples of hierarchical file mappings of an Optimal Flexible Architecture-compliant installation.

 $D:\E:\F:\$

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 B-2 Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
	System directory
C:\	
	Subtree for application software
C:\app\	
C:\app\oraInventory	Central Oralnventory 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.
C:\app\oracle\	Oracle base directory for user oracle. There can be many Oracle Database installations on a server, and many Oracle Database software installation owners.
	Oracle software homes that an Oracle installation owner owns should be located in the Oracle base directory for the Oracle software installation owner, unless that Oracle software is Oracle Grid Infrastructure deployed for a cluster.



Table B-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
C:\app\grid\username	Oracle base directory for the Oracle Grid Infrastructure software, where <i>username</i> is the name of the user that performed the software installation. 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.
C:\app\oracle\admin\	Subtree for database administration files
C:\app\oracle\admin\TAR	Subtree for support log files
C:\app\oracle\admin\db_sales\	admin subtree for database named "sales"
C:\app\oracle\admin\db_dwh\	admin subtree for database named "dwh"
C:\app\oracle\fast_recovery_are a\	Subtree for recovery files
C:\app\oracle\fast_recovery_are a\db_sales	Recovery files for database named "sales"
C:\app\oracle\fast_recovery_are a\db_dwh	Recovery files for database named "dwh"
D:\app\oracle\oradata E:\app\oracle\oradata F:\app\oracle\oradata	Oracle data file directories
C:\app\oracle\product\	Common path for Oracle software products other than Oracle Grid Infrastructure for a cluster.
C:\app\oracle\product\21.0.0\db home_1	Oracle home directory for Oracle Database 1, owned by Oracle Database installation owner account oracle.



Table B-2 (Cont.) Optimal Flexible Architecture Hierarchical File Path Examples

Directory	Description
C:\app\oracle\product\21.0.0\db home_2	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account oracle
C:\app\oradbowner\product\21. 0.0\dbhome_2	Oracle home directory for Oracle Database 2, owned by Oracle Database installation owner account oradbowner.
C:\app\oracle\product\21.0.0\gri	Oracle home directory for Oracle Grid Infrastructure for a standalone server, owned by Oracle Database and Oracle Grid Infrastructure installation owner oracle.
C:\app\21.0.0\grid	Oracle home directory for Oracle Grid Infrastructure for a cluster (Grid home), owned by user grid before installation, and owned by root after installation.
C:\app\grid\username\diag\crs\h ostname\crs\trace	Oracle Clusterware log files