

Oracle® Real Application Clusters

Installation and Configuration Guide

10g Release 1 (10.1) for AIX-Based Systems, hp HP-UX
PA-RISC (64-bit), hp Tru64 UNIX, Linux, Solaris Operating Sys-
tem (SPARC 64-bit)

Part No. B10766-01

January 2004

Oracle Real Application Clusters Installation and Configuration Guide 10g Release 1 (10.1) for AIX-Based Systems, hp HP-UX PA-RISC (64-bit), hp Tru64 UNIX, Linux, Solaris Operating System (SPARC 64-bit)

Part No. B10766-01

Copyright © 2004, Oracle. All rights reserved.

Primary Authors: David Austin, Mark Bauer, Kevin Flood, Emily Murphy

Contributing Authors: Jonathan Creighton, Raj Kumar

Contributors: Karin Brandauer, Jack Cai, Rajiv Jayaraman, Diana Lorentz, Randy Neville, Michael Polaski, Sudheendra Sampath, Clive Simpkins, Khethavath P. Singh

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Send Us Your Comments	xiii
Preface	xv
Intended Audience.....	xvi
Documentation Accessibility	xvi
Structure	xvi
Related Documents	xviii
Conventions	xix
What's New in Oracle Database 10g RAC Installation and Configuration?	xxiii
Oracle Database 10g Release 1 (10.1) New Features for RAC Installation and Configuration....	xxiv
Part I Oracle Database 10g Real Application Clusters Installation Planning and Requirements	
1 Introduction to Installing and Configuring Oracle Database 10g RAC	
General System Installation Requirements for Real Application Clusters	1-2
Hardware and Network Requirements for Oracle Database 10g Real Application Clusters.	1-2
Software Requirements for Oracle Database 10g Real Application Clusters	1-2
Cluster Setup and Pre-Installation Configuration Tasks for Real Application Clusters	1-3
Pre-Installation, Installation, and Post-Installation Overview	1-3
Pre-Installation for Oracle Database 10g Real Application Clusters	1-3
Installation for Oracle Database 10g Real Application Clusters	1-3
Post-Installation for Oracle Database 10g Real Application Clusters	1-4
The Oracle Universal Installer and Real Application Clusters	1-4
Storage Considerations for Installing Oracle Database 10g Real Application Clusters	1-5
Overview of Automatic Storage Management	1-5
Additional Considerations for Using Oracle Database 10g Features in RAC	1-7
Oracle Database 10g and Real Application Clusters	1-7
The Cluster Ready Services Clusterware.....	1-8
The Installed Real Application Clusters Components.....	1-8
Oracle Database 10g Real Application Clusters Version Compatibility	1-8
Required Groups	1-9

Part II Real Application Clusters Platform-Specific Pre-Installation Procedures

2 Pre-Installation Tasks for Installing RAC on AIX-Based Systems

Log In to the System as root	2-2
Check the Hardware Requirements	2-3
Check the Network Requirements	2-4
Network Hardware Requirements	2-4
IP Address Requirements	2-4
Checking the Network Requirements.....	2-5
Checking the Software Requirements	2-6
Check for Required Software	2-6
Check for Required Patches.....	2-7
Create Required UNIX Groups and User	2-9
Creating the Oracle Inventory Group	2-10
Determine Whether the Oracle Inventory Group Exists.....	2-10
Create the Oracle Inventory Group.....	2-11
Creating the OSDBA Group	2-11
Creating an OSOPER Group (Optional)	2-11
Creating the Oracle Software Owner User	2-12
Determine Whether an Existing Oracle Software Owner User Exists	2-12
Create a New Oracle Software Owner User	2-12
Modifying an Existing Oracle Software Owner User	2-13
Verifying That the UNIX User nobody Exists.....	2-13
Create Identical Users and Groups on Other Cluster Nodes	2-13
Identifying the User and Group IDs	2-14
Create the User and Groups on the Other Cluster Nodes	2-14
Set Up User Equivalence for rsh and rcp on All Cluster Nodes	2-15
Configure Shell Limits and Kernel Parameters	2-16
Configure Shell Limits.....	2-16
Configure System Configuration Parameters	2-16
Identify Required Software Directories	2-17
Oracle Base Directory	2-17
Oracle Inventory Directory	2-18
CRS Home Directory	2-18
Oracle Home Directory	2-18
Identify or Create an Oracle Base Directory	2-19
Identifying an Existing Oracle Base Directory	2-19
Creating a New Oracle Base Directory	2-21
Configure Disk Storage for Oracle CRS and Database Files	2-21
Choosing a Storage Option for Oracle CRS Files	2-22
Choosing a Storage Option for Oracle Database Files.....	2-22
Choosing a Storage Option for Oracle Database Recovery Files	2-22
Configuring Disk Storage	2-22
Create Directories for Oracle CRS, Database, or Recovery Files	2-23
Guidelines for Placing Oracle CRS Files on a File System	2-23
Guidelines for Placing Oracle Database Files on a File System.....	2-23
Guidelines for Placing Oracle Recovery Files on a File System	2-24

Creating Required Directories.....	2-25
Configure Disks for Automatic Storage Management	2-26
Identifying Storage Requirements for ASM.....	2-26
Using an Existing ASM Disk Group.....	2-29
Configuring Disks for ASM.....	2-30
Configure Raw Logical Volumes.....	2-31
Review Important Information	2-31
Create a Volume Group	2-32
Create Raw Logical Volumes in the New Volume Group	2-34
Import the Volume Group on the Other Cluster Nodes	2-36
Activate the Volume Group in Concurrent Mode on All Cluster Nodes	2-37
Create the DBCA Raw Device Mapping File	2-37
Verifying the Cluster Software Configuration	2-39
Stop Existing Oracle Processes	2-40
Configure the oracle User's Environment	2-41

3 Pre-Installation Tasks for Installing RAC on HP-UX Systems

Log In to the System as root	3-2
Check the Hardware Requirements.....	3-3
Check the Network Requirements.....	3-4
Network Hardware Requirements	3-4
IP Address Requirements	3-4
Checking the Network Requirements.....	3-5
Check the Software Requirements	3-6
Check for Required Software	3-6
Check for Required Patches.....	3-7
Create Required Symbolic Links.....	3-9
Create Required UNIX Groups and User	3-9
Creating the Oracle Inventory Group	3-11
Determine Whether the Oracle Inventory Group Exists.....	3-11
Create the Oracle Inventory Group.....	3-11
Creating the OSDBA Group	3-11
Creating an OSOPER Group (Optional)	3-12
Creating the Oracle Software Owner User	3-12
Determine Whether an Existing Oracle Software Owner User Exists	3-12
Create a New Oracle Software Owner User	3-13
Modify an Existing Oracle Software Owner User.....	3-13
Creating an Unprivileged User	3-13
Create Identical Users and Groups on Other Cluster Nodes	3-13
Identifying the User and Group IDs	3-13
Create the User and Groups on the Other Cluster Nodes	3-14
Set Up User Equivalence for rsh and rcp on All Cluster Nodes	3-15
Grant Privileges to the OSDBA Group	3-15
Configure Kernel Parameters.....	3-16
Identify Required Software Directories	3-18
Oracle Base Directory	3-18
Oracle Inventory Directory	3-19

CRS Home Directory	3-19
Oracle Home Directory	3-19
Identify or Create an Oracle Base Directory	3-20
Identifying an Existing Oracle Base Directory	3-20
Creating a New Oracle Base Directory	3-21
Configure Disk Storage for Oracle CRS and Database Files.....	3-22
Choosing a Storage Option for Oracle CRS Files	3-22
Choosing a Storage Option for Oracle Database Files.....	3-22
Planning a Storage Option for Oracle Database Recovery Files	3-22
Configuring Disk Storage	3-23
Configure Disks for Automatic Storage Management	3-23
Identifying Storage Requirements for ASM.....	3-23
Using an Existing ASM Disk Group.....	3-25
Configuring Disks for ASM.....	3-27
Configure Raw Logical Volumes.....	3-27
Review Important Information	3-28
Disable Operating System Activation of Shared Volume Groups.....	3-29
Create a Volume Group	3-30
Create Raw Logical Volumes in the New Volume Group	3-31
Export the Volume Group and Import It on the Other Cluster Nodes.....	3-33
Activate the Volume Group in Shared Mode on All Cluster Nodes	3-34
Create the DBCA Raw Device Mapping File	3-34
Verify the Cluster Software Configuration.....	3-36
Stop Existing Oracle Processes	3-36
Configure the oracle User's Environment	3-37

4 Pre-Installation Tasks for Installing RAC on HP Tru64 UNIX

Log In to the System as root	4-2
Check the Hardware Requirements.....	4-3
Check the Network Requirements.....	4-4
Network Hardware Requirements	4-4
IP Address Requirements	4-4
Checking the Network Requirements.....	4-5
Check the Software Requirements	4-5
Check for Required Patches.....	4-7
Create Required UNIX Groups and User	4-8
Creating the Oracle Inventory Group	4-10
Creating the OSDBA Group	4-10
Creating an OSOPER Group (Optional)	4-10
Creating the Oracle Software Owner User	4-11
Determine Whether an Existing Oracle Software Owner User Exists	4-11
Create a New Oracle Software Owner User	4-11
Modify an Existing Oracle Software Owner User.....	4-12
Verifying That the UNIX User nobody Exists.....	4-12
Create Identical Users and Groups on Other Cluster Nodes	4-12
Identifying the User and Group IDs	4-12
Create the User and Groups on the Other Cluster Nodes	4-13

Set Up User Equivalence for rsh and rcp on All Cluster Nodes	4-14
Configure Kernel Subsystem Attributes	4-15
Identify Required Software Directories	4-16
Oracle Base Directory	4-17
Oracle Inventory Directory	4-17
CRS Home Directory	4-18
Oracle Home Directory	4-18
Identify or Create an Oracle Base Directory	4-18
Identifying an Existing Oracle Base Directory	4-19
Creating a New Oracle Base Directory	4-20
Configure Disk Storage for Oracle CRS and Database Files	4-21
Choosing a Storage Option for Oracle Database Files	4-21
Choosing a Storage Option for Oracle Database Recovery Files	4-22
Configuring Disk Storage	4-22
Create Directories for Oracle CRS, Database, or Recovery Files	4-22
Configure Disks for Automatic Storage Management	4-25
Identifying Storage Requirements for ASM.....	4-26
Using an Existing ASM Disk Group.....	4-28
Configuring Disks for ASM.....	4-29
Verify that the Cluster Software is Running	4-31
Stop Existing Oracle Processes	4-31
Configure the oracle User's Environment	4-32

5 Pre-Installation Tasks for Installing RAC on Linux-Based Systems

Log In to the System as root	5-2
Check the Hardware Requirements	5-3
Check the Network Requirements	5-4
Network Hardware Requirements	5-4
IP Address Requirements	5-4
Checking the Network Requirements.....	5-5
Check the Software Requirements	5-5
Check for Required Software	5-6
Check for Required Patches.....	5-8
Create Required UNIX Groups and User	5-9
Creating the Oracle Inventory Group	5-10
Determine Whether the Oracle Inventory Group Exists.....	5-10
Create the Oracle Inventory Group	5-10
Creating the OSDBA Group	5-11
Creating an OSOPER Group (Optional)	5-11
Creating the Oracle Software Owner User	5-11
Determine Whether an Existing Oracle Software Owner User Exists	5-11
Create a New Oracle Software Owner User	5-12
Modify an Existing Oracle Software Owner User.....	5-12
Verifying That the UNIX User nobody Exists.....	5-12
Create Identical Users and Groups on Other Cluster Nodes	5-13
Identifying the User and Group IDs	5-13
Create the User and Groups on the Other Cluster Nodes	5-13

Configure SSH on All Cluster Nodes.....	5-14
Configure Kernel Parameters and Shell Limits	5-16
Configuring Kernel Parameters	5-16
Set Shell Limits for the oracle User.....	5-18
Identify Required Software Directories	5-18
Oracle Base Directory	5-19
Oracle Inventory Directory	5-19
CRS Home Directory	5-20
Oracle Home Directory	5-20
Identify or Create an Oracle Base Directory	5-20
Identifying an Existing Oracle Base Directory	5-20
Creating a New Oracle Base Directory	5-22
Configure Disk Storage for Oracle CRS and Database Files.....	5-23
Choosing a Storage Option for Oracle CRS Files	5-23
Choosing a Storage Option for Oracle Database Files.....	5-23
Choosing a Storage Option for Oracle Database Recovery Files	5-23
Configuring Disk Storage	5-24
Create Directories for Oracle CRS, Database, or Recovery Files	5-24
Guidelines for Placing Oracle CRS Files on a File System	5-24
Guidelines for Placing Oracle Database Files on a File System.....	5-24
Guidelines for Placing Oracle Recovery Files on a File System	5-25
Creating Required Directories.....	5-26
Configure Disks for Automatic Storage Management	5-27
Identifying Storage Requirements for ASM.....	5-27
Using an Existing ASM Disk Group.....	5-30
Configuring Disks for ASM.....	5-31
Configuring Disks for ASM Using the ASM Library Driver	5-31
Configuring Disks for ASM Using Raw Devices	5-34
Configure Raw Partitions	5-36
Review Important Information	5-37
Create Raw Partitions	5-37
Bind the Partitions to Raw Devices	5-39
Create the DBCA Raw Device Mapping File	5-40
Verify that the Required Software is Running.....	5-41
Check the Configuration of the hangcheck-timer Module	5-42
Check the Configuration of OCFS	5-42
Stop Existing Oracle Processes	5-43
Configure the oracle User's Environment	5-44

6 Pre-Installation Tasks for Installing RAC on Solaris Operating System (SPARC 64-Bit)

Log In to the System as root	6-2
Check the Hardware Requirements.....	6-3
Check the Network Requirements.....	6-4
Network Hardware Requirements	6-4
IP Address Requirements	6-5
Checking the Network Requirements.....	6-6

Check the Software Requirements	6-6
Check for Required Software	6-6
Check for Required Patches.....	6-8
Create Required UNIX Groups and User	6-10
Creating the Oracle Inventory Group	6-11
Determine Whether the Oracle Inventory Group Exists.....	6-12
Create the Oracle Inventory Group.....	6-12
Creating the OSDBA Group	6-12
Creating an OSOPER Group (Optional)	6-12
Creating the Oracle Software Owner User	6-12
Determine Whether an Existing Oracle Software Owner User Exists	6-13
Create a New Oracle Software Owner User	6-13
Modify an Existing Oracle Software Owner User.....	6-14
Verifying That the UNIX User nobody Exists.....	6-14
Create Identical Users and Groups on Other Cluster Nodes	6-14
Identifying the User and Group IDs	6-14
Create the User and Groups on the Other Cluster Nodes	6-14
Configure SSH on All Cluster Nodes.....	6-15
Set Up User Equivalence for rsh and rcp on All Cluster Nodes	6-17
Configure Kernel Parameters.....	6-18
Identify Required Software Directories	6-19
Oracle Base Directory	6-19
Oracle Inventory Directory	6-20
CRS Home Directory	6-21
Oracle Home Directory	6-21
Identify or Create an Oracle Base Directory	6-21
Identifying an Existing Oracle Base Directory	6-21
Creating a New Oracle Base Directory	6-23
Configure Disk Storage for Oracle CRS and Database Files.....	6-24
Choosing a Storage Option for Oracle CRS Files	6-24
Special Considerations for Sun Cluster: Using VERITAS Volume Manager.....	6-25
Sun Cluster Private Host Names	6-25
Choosing a Storage Option for Oracle Database Recovery Files	6-25
Configuring Disk Storage	6-25
Create Directories for Oracle CRS, Database, or Recovery Files	6-26
Guidelines for Placing Oracle CRS Files on a File System	6-26
Guidelines for Placing Oracle Database Files on a File System.....	6-26
Guidelines for Placing Oracle Recovery Files on a File System	6-27
Creating Required Directories.....	6-28
Configure Disks for Automatic Storage Management	6-29
Identifying Storage Requirements for ASM.....	6-30
Using an Existing ASM Disk Group.....	6-32
Configuring Disks for ASM.....	6-33
Configure Raw Logical Volumes.....	6-34
Create a Shared Disk Group	6-35
Create Raw Logical Volumes in the New Disk Group	6-36
Deploy the Disk Group and Import It on the Other Cluster Nodes	6-38

Create the DBCA Raw Device Mapping File	6-38
Verify that the Cluster Software is Configured and Running.....	6-40
Stop Existing Oracle Processes	6-41
Configure the oracle User's Environment	6-42

Part III Installing CRS and Oracle Database 10g with RAC, Creating RAC Databases, and Performing Post-Installation Tasks

7 Installing Cluster Ready Services

Installation Setup Procedures.....	7-2
Installing the UDLM for SunClusters.....	7-2
Installing Cluster Ready Services with the OUI	7-2
Cluster Ready Services Background Processes.....	7-5

8 Installing Oracle Database 10g with Real Application Clusters

Selecting a Database Configuration Type	8-2
Configuration Type Descriptions	8-2
General Purpose, Transaction Processing, and Data Warehouse Configuration Types ..	8-2
Using The Customized Configuration Type.....	8-3
Processing of the OUI, the DBCA, and Other Assistants During Installation	8-3
Installation Setup Procedures.....	8-4
Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer	8-4
De-Installing Real Application Clusters Software.....	8-8
De-Installing Oracle Database 10g RAC Software	8-8
De-Installing Cluster Ready Services	8-9

9 Creating RAC Databases with the Database Configuration Assistant

Using the Database Configuration Assistant in Real Application Clusters.....	9-2
Benefits of Using the Database Configuration Assistant.....	9-2
Real Application Clusters High Availability Services.....	9-2
Service Configuration and Instance Preferences	9-2
Transparent Application Failover Policies	9-3
Creating the Database after Installation Using the Database Configuration Assistant	9-3
Creating a Real Application Clusters Database with the DBCA	9-3
Deleting a Real Application Clusters Database with the DBCA.....	9-9

10 Real Application Clusters Post-Installation Procedures

Required Post-Installation Tasks	10-2
Download and Install Patches.....	10-2
Configure Oracle Products	10-2
Recommended Post-Installation Tasks	10-3
Back Up the root.sh Script.....	10-3
Set Up Users Accounts	10-3
Verifying Enterprise Manager Operations	10-3

Part IV Real Application Clusters Environment Configuration

11 Configuring the Server Parameter File in Real Application Clusters Environments

Parameter Files and Real Application Clusters.....	11-2
Using Server Parameter Files in Real Application Clusters	11-2
Location of The Server Parameter File	11-2
Parameter File Search Order in Real Application Clusters.....	11-3
Migrating to the Server Parameter File in Real Application Clusters Environments	11-3
Server Parameter File Placement in Real Application Clusters	11-3
Procedures for Migrating to the Server Parameter File	11-3
Server Parameter File Errors in Real Application Clusters.....	11-4

12 Understanding the Real Application Clusters Installed Configuration

Understanding the Configured Environment in Real Application Clusters.....	12-2
The Oracle Cluster Registry in Real Application Clusters	12-2
oratab Configurations for Real Application Clusters	12-2
Database Components Created Using the Database Configuration Assistant	12-2
Tablespaces and Datafiles	12-2
Control Files	12-3
Redo Log Files	12-3
Managing Undo Tablespaces in Real Application Clusters	12-3
Initialization Parameter Files	12-4
Configuring Service Registration-Related Parameters in Real Application Clusters	12-4
Configuring the Listener File (listener.ora).....	12-5
Local Listeners	12-5
Multiple Listeners	12-5
How Oracle Uses the Listener (listener.ora File).....	12-5
Listener Registration and PMON Discovery	12-6
Directory Server Access (ldap.ora File)	12-6
Net Service Names (tnsnames.ora File).....	12-7
Profile (sqlnet.ora File).....	12-11

Part V Real Application Clusters Installation and Configuration Reference Information

A Troubleshooting the Real Application Clusters Installation Process

Troubleshooting the Real Application Clusters Installation.....	A-2
Real Application Clusters Installation Error Messages	A-2
Performing Cluster Diagnostics During Real Application Clusters Installations	A-2

B Using Scripts to Create Real Application Clusters Databases

Creating a Database Using Scripts.....	B-2
--	-----

C Configuring Raw Devices for Real Application Clusters

Configuring Shared Disk Subsystems for Real Application Clusters	C-2
Additional Raw Devices Required by the DBCA for Non-CFS Environments	C-2
Planning Your Raw Device Creation Strategy	C-2

D Converting to Real Application Clusters from Single-Instance Oracle Databases

Deciding to Convert	D-2
Prerequisites for Conversion	D-2
Single-Instance to Cluster-Enabled Conversion Administrative Issues	D-2
Converting from Single-Instance to Real Application Clusters	D-2
Single Instance on a Non-Cluster Machine to Oracle Database 10g with RAC	D-2
Back up the Original Single-Instance Database	D-3
Perform the Pre-Installation Steps.....	D-3
Set up the Cluster	D-3
Copy the Preconfigured Database Image.....	D-3
Install Oracle Database 10g Software with Real Application Clusters	D-3
Single Instance on a Cluster to Oracle Database 10g RAC.....	D-4
Single Instance on a Cluster Running from a Cluster Enabled Oracle Home	D-4
Automated Conversion Procedure.....	D-5
Manual Conversion Procedure	D-5
Single Instance on a Cluster Running from a RAC-Disabled Oracle Home	D-6
Single Instance on a Cluster Running from non-Cluster Installed Oracle Home	D-7
Post-Conversion Steps	D-7

E Directory Structure for Oracle Database 10g Real Application Clusters Environments

Understanding the Real Application Clusters Directory Structure	E-2
UNIX Directory Structures for Real Application Clusters	E-2

Index

Send Us Your Comments

Oracle Real Application Clusters Installation and Configuration Guide 10g Release 1 (10.1) for AIX-Based Systems, hp HP-UX PA-RISC (64-bit), hp Tru64 UNIX, Linux, Solaris Operating System (SPARC 64-bit)

Part No. B10766-01

Oracle welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the title and part number of the documentation and the chapter, section, and page number (if available). You can send comments to us in the following ways:

- Electronic mail: infodev_us@oracle.com
- FAX: (650) 506-7227. Attn: Server Technologies Documentation Manager
- Postal service:

Oracle Corporation
Server Technologies Documentation Manager
Street Address, Mailstop 4op11
Redwood Shores, CA 94065
USA

If you would like a reply, please give your name, address, telephone number, and electronic mail address (optional).

If you have problems with the software, please contact your local Oracle Support Services.

Preface

The *Oracle Real Application Clusters Installation and Configuration Guide* explains how to install and configure Oracle Real Application Clusters (RAC). With the exception of the platform-specific sections, information in this manual applies to Oracle Database 10g RAC as it runs on most operating systems. This preface contains the following topics:

- [Intended Audience](#)
- [Documentation Accessibility](#)
- [Structure](#)
- [Related Documents](#)
- [Conventions](#)

See Also: *Oracle Database System Administration Guide 10g Release 1 (10.1) for IBM z/OS (OS/390)* for more information about installing RAC on the IBM z/OS platform

Note: Some platforms, such as hp Itanium64, MAC OS X, and Solaris x86, are not included in this document. Refer to Oracle Database 10g release notes for additional platform information.

Intended Audience

The *Oracle Real Application Clusters Installation and Configuration Guide* is primarily for network or Database Administrators (DBAs) who install and configure RAC.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For additional information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation JAWS, a Windows screen reader, may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, JAWS may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Structure

This document contains the following five parts:

Part I: "Oracle Database 10g Real Application Clusters Installation Planning and Requirements"

Part I introduces the RAC installation process.

Chapter 1, "Introduction to Installing and Configuring Oracle Database 10g RAC"

This chapter describes the RAC installation process and provides RAC installation planning information.

Part II: Real Application Clusters Platform-Specific Pre-Installation Procedures

Part II describes the platform-specific pre-installation procedures for installing RAC.

Chapter 2, "Pre-Installation Tasks for Installing RAC on AIX-Based Systems"

This chapter describes the pre-installation procedures for installing RAC on IBM AIX systems.

Chapter 3, "Pre-Installation Tasks for Installing RAC on HP-UX Systems"

This chapter describes the pre-installation procedures for installing RAC on HP-UX systems.

Chapter 4, "Pre-Installation Tasks for Installing RAC on HP Tru64 UNIX"

This chapter describes the pre-installation procedures for installing RAC on HP Tru64 UNIX systems.

Chapter 5, "Pre-Installation Tasks for Installing RAC on Linux-Based Systems"

This chapter describes the pre-installation procedures for installing RAC on Linux-based systems.

Chapter 6, "Pre-Installation Tasks for Installing RAC on Solaris Operating System (SPARC 64-Bit)"

This chapter describes the pre-installation procedures for installing RAC on Solaris Operating System (SPARC 64-Bit) systems.

Part III: "Installing CRS and Oracle Database 10g with RAC, Creating RAC Databases, and Performing Post-Installation Tasks"

Part III describes how to install Cluster Ready Services and Oracle Database 10g with Real Application Clusters.

Chapter 7, "Installing Cluster Ready Services"

This chapter describes the pre-installation procedures for installing Cluster Ready Services.

Chapter 8, "Installing Oracle Database 10g with Real Application Clusters"

This chapter describes how to install Oracle Database 10g with Real Application Clusters on all operating systems.

Chapter 9, "Creating RAC Databases with the Database Configuration Assistant"

This chapter explains how to use the Database Configuration Assistant to create RAC databases.

Chapter 10, "Real Application Clusters Post-Installation Procedures"

This chapter describes the post-installation tasks for RAC.

Part IV: Real Application Clusters Environment Configuration

Part IV provides Oracle Database 10g Real Application Clusters environment configuration information.

Chapter 11, "Configuring the Server Parameter File in Real Application Clusters Environments"

This chapter describes the use of the server parameter file (SPFILE) in Real Application Clusters.

Chapter 12, "Understanding the Real Application Clusters Installed Configuration"

This chapter describes the Oracle Database 10g Real Application Clusters installed configuration.

Part V: Real Application Clusters Installation and Configuration Reference Information

Part V provides reference information for the installation and configuration of RAC.

Appendix A, "Troubleshooting the Real Application Clusters Installation Process"

This appendix provides RAC installation and configuration troubleshooting information.

Appendix B, "Using Scripts to Create Real Application Clusters Databases"

This appendix explains how to use scripts in RAC.

Appendix C, "Configuring Raw Devices for Real Application Clusters"

This appendix explains how to configure shared disk subsystems using raw devices in RAC environments.

Appendix D, "Converting to Real Application Clusters from Single-Instance Oracle Databases"

This appendix describes how to convert to Oracle Database 10g RAC from single-instance Oracle databases.

Appendix E, "Directory Structure for Oracle Database 10g Real Application Clusters Environments"

This appendix describes the directory structure for the installed RAC software on UNIX-based systems.

Related Documents

For more information, refer to these Oracle resources:

- *Oracle Real Application Clusters Administrator's Guide*
- *Oracle Real Application Clusters Deployment and Performance Guide*

Error messages are only available online or by using Tahiti, the Oracle documentation search tool.

Installation Guides

- *Oracle Diagnostics Pack Installation*

Note: Some platforms, such as hp Itanium64, MAC OS X, and Solaris x86, are not included in this document. Refer to Oracle Database 10g release notes for additional platform information.

Operating System-Specific Administrative Guides

- *Oracle Database Administrator's Reference, 10g Release 1 (10.1) for UNIX Systems*
- *Oracle Database System Administration Guide 10g Release 1 (10.1) for IBM z/OS (OS/390)*

Oracle Database 10g Real Application Clusters Management

- *Oracle Real Application Clusters Administrator's Guide*
- *Oracle 2 Day DBA*
- *Getting Started with the Oracle Diagnostics Pack*

Generic Documentation

- *Oracle Database New Features*

- *Oracle Database Concepts*
- *Oracle Net Services Administrator's Guide*
- *Oracle Database Reference*

Printed documentation is available for sale in the Oracle Store at

<http://oraclestore.oracle.com/>

To download free release notes, installation documentation, white papers, or other collateral, please visit the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at

<http://otn.oracle.com/membership/>

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at

<http://otn.oracle.com/documentation/>

Oracle error message documentation is only available in HTML. If you only have access to the Oracle Documentation CD, then browse the error messages by range. Once you find a range, use your browser's "find in page" feature to locate a specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

Printed documentation is available for sale in the Oracle Store at

<http://oraclestore.oracle.com/>

To download free release notes, installation documentation, white papers, or other collateral, please visit the Oracle Technology Network (OTN). You must register online before using OTN; registration is free and can be done at

<http://otn.oracle.com/membership/>

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at

<http://otn.oracle.com/documentation/>

Conventions

This section describes the conventions used in the text and code examples of this documentation set. It describes:

- [Conventions in Text](#)
- [Conventions in Code Examples](#)

Conventions in Text

We use various conventions in text to help you more quickly identify special terms. The following table describes those conventions and provides examples of their use.

Convention	Meaning	Example
Bold	Bold typeface indicates terms that are defined in the text or terms that appear in a glossary, or both.	When you specify this clause, you create an index-organized table .

Convention	Meaning	Example
<i>Italics</i>	Italic typeface indicates book titles or emphasis.	<i>Oracle Database Concepts</i> Ensure that the recovery catalog and target database do <i>not</i> reside on the same disk.
UPPERCASE monospace (fixed-width) font	Uppercase monospace typeface indicates elements supplied by the system. Such elements include parameters, privileges, datatypes, RMAN keywords, SQL keywords, SQL*Plus or utility commands, packages and methods, as well as system-supplied column names, database objects and structures, usernames, and roles.	You can specify this clause only for a NUMBER column. You can back up the database by using the BACKUP command. Query the TABLE_NAME column in the USER_TABLES data dictionary view. Use the DBMS_STATS.GENERATE_STATS procedure.
lowercase monospace (fixed-width) font	Lowercase monospace typeface indicates executables, filenames, directory names, and sample user-supplied elements. Such elements include computer and database names, net service names, and connect identifiers, as well as user-supplied database objects and structures, column names, packages and classes, usernames and roles, program units, and parameter values. Note: Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	Enter sqlplus to start SQL*Plus. The password is specified in the orapwd file. Back up the datafiles and control files in the /disk1/oracle/dbs directory. The department_id, department_name, and location_id columns are in the hr.departments table. Set the QUERY_REWRITE_ENABLED initialization parameter to true. Connect as oe user. The JRepUtil class implements these methods.
<i>lowercase italic monospace (fixed-width) font</i>	Lowercase italic monospace font represents placeholders or variables.	You can specify the <i>parallel_clause</i> . Run <i>old_release</i> .SQL where <i>old_release</i> refers to the release you installed prior to upgrading.

Conventions in Code Examples

Code examples illustrate SQL, PL/SQL, SQL*Plus, or other command-line statements. They are displayed in a monospace (fixed-width) font and separated from normal text as shown in this example:

```
SELECT username FROM dba_users WHERE username = 'MIGRATE';
```

The following table describes typographic conventions used in code examples and provides examples of their use.

Convention	Meaning	Example
[]	Brackets enclose one or more optional items. Do not enter the brackets.	DECIMAL (<i>digits</i> [, <i>precision</i>])
{ }	Braces enclose two or more items, one of which is required. Do not enter the braces.	{ENABLE DISABLE}
	A vertical bar represents a choice of two or more options within brackets or braces. Enter one of the options. Do not enter the vertical bar.	{ENABLE DISABLE} [COMPRESS NOCOMPRESS]

Convention	Meaning	Example
...	Horizontal ellipsis points indicate either: <ul style="list-style-type: none"> ■ That we have omitted parts of the code that are not directly related to the example ■ That you can repeat a portion of the code 	<pre>CREATE TABLE ... AS subquery; SELECT col1, col2, ... , coln FROM employees;</pre>
. . .	Vertical ellipsis points indicate that we have omitted several lines of code not directly related to the example.	<pre>SQL> SELECT NAME FROM V\$DATAFILE; NAME ----- /fs1/dbs/tbs_01.dbf /fs1/dbs/tbs_02.dbf . . . /fs1/dbs/tbs_09.dbf 9 rows selected.</pre>
Other notation	You must enter symbols other than brackets, braces, vertical bars, and ellipsis points as shown.	<pre>acctbal NUMBER(11,2); acct CONSTANT NUMBER(4) := 3;</pre>
<i>Italics</i>	Italicized text indicates placeholders or variables for which you must supply particular values.	<pre>CONNECT SYSTEM/system_password DB_NAME = database_name</pre>
UPPERCASE	Uppercase typeface indicates elements supplied by the system. We show these terms in uppercase in order to distinguish them from terms you define. Unless terms appear in brackets, enter them in the order and with the spelling shown. However, because these terms are not case sensitive, you can enter them in lowercase.	<pre>SELECT last_name, employee_id FROM employees; SELECT * FROM USER_TABLES; DROP TABLE hr.employees;</pre>
lowercase	Lowercase typeface indicates programmatic elements that you supply. For example, lowercase indicates names of tables, columns, or files. Note: Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	<pre>SELECT last_name, employee_id FROM employees; sqlplus hr/hr CREATE USER mjones IDENTIFIED BY ty3MU9;</pre>

What's New in Oracle Database 10g RAC Installation and Configuration?

This section describes the Oracle Database 10g release 1 (10.1) features as they pertain to the installation and configuration of Real Application Clusters (RAC). The topic in this section is:

- [Oracle Database 10g Release 1 \(10.1\) New Features for RAC Installation and Configuration](#)

Oracle Database 10g Release 1 (10.1) New Features for RAC Installation and Configuration

- This book contains Oracle Database 10g pre-installation and installation instructions for UNIX-based platforms on which RAC operates.
- The Oracle Database 10g with RAC is available on both the Standard Edition and the Enterprise Edition.
- The Oracle Database 10g installation requires you to perform a two-phase process in which you run the Oracle Universal Installer (OUI) twice. The first phase installs Oracle Cluster Ready Services Release 1 (10.1.0.2) and the second phase installs the Oracle Database 10g software with RAC. The installation also enables you to create and configure services for your RAC environment. If you have a previous Oracle cluster database version, then the OUI activates the Database Upgrade Assistant (DBUA) to automatically upgrade your pre-Oracle Database 10g cluster database. The Oracle Database 10g installation process provides single system image, ease of use, and accuracy for RAC installations and patches.
- Cluster Ready Services (CRS) provides Oracle Database 10g RAC high availability components that provide many system management features. CRS components also interact with the vendor clusterware, if present, to coordinate cluster membership information.
- There are new and changed pages and dialogs for the Oracle Universal Installer (OUI), the Database Configuration Assistant (DBCA), and the Database Upgrade Assistant. The Virtual Internet Protocol Configuration Assistant (VIPCA) is a new tool for this release. These enhancements are described in the following:
 - OUI Cluster Installation Mode Page—This page enables you to select whether to perform a cluster or a single-instance Oracle Database 10g installation.
 - SYS and SYSTEM Passwords Page—This page has fields for entering and confirming the SYS and SYSTEM user passwords. This includes SYSMAN and DBSNMP if you use Oracle Enterprise Manager Database Control.
 - Storage Options Page—This page has storage options for selecting the storage type for the database files such as control files, datafiles, and redo logs.
 - DBCA Services Page—This page enables you to create and configure services for your RAC environment.
 - DBCA Initialization Parameters Page—This page has two dialogs to display both Basic and Advanced parameter settings.
 - VIPCA—The pages for this assistant enable you to configure virtual internet protocol addresses for your RAC database.
- A new auxiliary, system-managed tablespace called SYSAUX contains performance data and combines content that was stored in different tablespaces (some of which are no longer required) in earlier releases. This is a required tablespace for which you must plan disk space.
- The `gsdctl` commands are obsolete. The CRS installation process stops any existing GSD processes.
- Pre-Oracle Database 10g cluster manager implementations on some platforms were referred to as "Cluster Manager". The cluster manager on all platforms in Oracle Database 10g is known as Cluster Synchronization Services (CSS). The Oracle Cluster Synchronization Service Daemon (OCSSD) performs this function.
- Oracle Database 10g provides cluster file system support for Linux.

See Also: <http://otn.oracle.com> for more information about Oracle Cluster File System on Linux

- RAC and the DBCA support Automatic Storage Management (ASM) and Oracle Managed Files (OMF).

See Also:

- The *Oracle Real Application Clusters Administrator's Guide* for more information about administering services and storage in RAC
 - The *Oracle Database Upgrade Guide* for information about using the DBUA
- The Oracle Database 10g version of the `srvConfig.loc` file is the `ocr.loc` file. The Oracle9i version of `srvConfig.loc` still exists for backward compatibility.

Part I

Oracle Database 10g Real Application Clusters Installation Planning and Requirements

Part I describes how to plan your Real Application Clusters (RAC) installation and describes the RAC installation requirements. The chapter in Part I is:

[Chapter 1, "Introduction to Installing and Configuring Oracle Database 10g RAC"](#)

Introduction to Installing and Configuring Oracle Database 10g RAC

This chapter provides an overview of the Real Application Clusters (RAC) installation and configuration procedures. This chapter includes the following topics:

- [General System Installation Requirements for Real Application Clusters](#)
- [Cluster Setup and Pre-Installation Configuration Tasks for Real Application Clusters](#)
- [Pre-Installation, Installation, and Post-Installation Overview](#)
- [The Oracle Universal Installer and Real Application Clusters](#)
- [Storage Considerations for Installing Oracle Database 10g Real Application Clusters](#)
- [Additional Considerations for Using Oracle Database 10g Features in RAC](#)
- [Oracle Database 10g Real Application Clusters Version Compatibility](#)
- [Oracle Database 10g and Real Application Clusters](#)

General System Installation Requirements for Real Application Clusters

Each node that is going to be part of your RAC installation must meet the following hardware and software requirements. You will also perform step-by-step tasks for hardware and software verification for the platform-specific pre-installation procedures in [Part II](#) of this book.

Before using this manual, however, you should read the *Oracle Real Application Clusters Deployment and Performance Guide* and the *Oracle Real Application Clusters Administrator's Guide*.

Hardware and Network Requirements for Oracle Database 10g Real Application Clusters

Each node in a cluster requires the following hardware:

- External shared disks.
- One private internet protocol (IP) address for each node to serve as the private interconnect. This IP address must be separate from the public network and it must have the same interface name on every node that is part of your cluster.
- One public IP address for each node to serve as the Virtual IP address for client connections and for connection failover. This is in addition to the operating-system managed public host IP address that is already assigned to the node by the operating system. This public Virtual IP must be associated with the same interface name on every node that is part of your cluster. In addition, the IP addresses that you use for all of the nodes that are part of a cluster must be from the same subnet. The host names for the VIP must be registered with the domain name server (DNS). The Virtual IP address should not be in use at the time of the installation because this is a Virtual IP address that Oracle manages.
- Redundant switches as a standard configuration for all cluster sizes.

Sections in each platform-specific pre-installation chapter that appear under the headings "Choosing a Storage Option for Oracle Database Files" and "Choosing a Storage Option for Oracle Database Recovery Files" describe the disk configuration options that are available for each platform. Review these options before you decide which storage option to use in your RAC environment.

Software Requirements for Oracle Database 10g Real Application Clusters

Each node in a cluster requires a supported interconnect software protocol, to support Cache Fusion, and Transmission Control Protocol/Internet Protocol (TCP/IP) to support Cluster Ready Services (CRS) polling. In addition to User Datagram Protocol (UDP), other supported vendor-specific interconnect protocols include Remote Shared Memory, Hyper Messaging protocol, and Reliable Data Gram. Note that Token Ring is not supported for cluster interconnects on AIX. Your interconnect *must* be certified by Oracle for your platform. You should also have a Web browser to view online documentation.

For functionality required from the vendor clusterware, Oracle's clusterware provides the equivalent functionality. Also, using Oracle clusterware reduces installation and support complications. However, vendor clusterware may be needed if customers use non-ethernet interconnect or if you have deployed clusterware-dependent applications on the same cluster where you deploy RAC.

Cluster Setup and Pre-Installation Configuration Tasks for Real Application Clusters

Before installing RAC, perform the following procedures:

1. Ensure that you have a certified combination of operating system and Oracle software version by referring to the Oracle *MetaLink* certification information under "Availability & Certification" > "1.View Certifications by Product" at the following site:

<http://metalink.oracle.com>

Note: The layout of the MetaLink site and the site's certification policies are subject to change.

2. Configure a high-speed interconnect that uses a private network. Configure a second interconnect for redundancy to avoid making the interconnect a single point of failure.
3. Determine the storage option for your system and configure the shared disk. Oracle recommends that you use Automatic Storage Management (ASM) and Oracle Managed Files (OMF), or a cluster file system. If you use ASM or a cluster file system, then you can also take advantage of OMF and other Oracle Database 10g storage features.
4. Install the operating system patches that are listed in the pre-installation chapters in this book in [Part II](#).

Pre-Installation, Installation, and Post-Installation Overview

The following describes the installation procedures that are described in [Part II](#) and [Part III](#) of this book.

Pre-Installation for Oracle Database 10g Real Application Clusters

The platform-specific pre-installation procedures in [Part II](#) explain how to verify user equivalence, perform network connectivity tests, as well as how to set directory and file permissions. Complete all of the pre-installation procedures and verify that your system meets all of the pre-installation requirements before proceeding to the install phase.

Installation for Oracle Database 10g Real Application Clusters

The Oracle Database 10g Real Application Clusters installation is a two-phase installation. In phase one, use the Oracle Universal Installer (OUI) to install CRS as described in [Chapter 7, "Installing Cluster Ready Services"](#). Note that the Oracle home that you use in phase one is a home for the CRS software which *must be different* from the Oracle home that you use in phase two for the installation of the Oracle database software with RAC components. The CRS pre-installation starts the CRS processes in preparation for installing Oracle Database 10g with RAC as described in [Chapter 8, "Installing Oracle Database 10g with Real Application Clusters"](#). You use the OUI in this phase to install the RAC software. [Chapter 8, "Installing Oracle Database 10g with Real Application Clusters"](#). The procedures in [Chapter 8](#) are generic for all platforms.

If the OUI detects Oracle cluster software from a previous release, then the OUI starts the Database Upgrade Assistant (DBUA) to upgrade your database to Oracle Database 10g release 1 (10.1). In addition, the DBUA displays a Service Configuration page for configuring services in your RAC database.

After the installation completes, the OUI starts the Oracle assistants, such as the Database Configuration Assistant (DBCA), to configure your environment and create your RAC database. You can later use the DBCA Instance Management feature to add or modify services and instances as described in [Chapter 9, "Creating RAC Databases with the Database Configuration Assistant"](#).

Post-Installation for Oracle Database 10g Real Application Clusters

After you create your database, download and install the most recent patch sets for your Oracle Database 10g version as described in [Chapter 10, "Real Application Clusters Post-Installation Procedures"](#). If you are using other Oracle products with your RAC database, then you must also configure them.

You must also perform several post-installation configuration tasks to use certain Oracle Database 10g products such as the Sample Schema, Oracle Net Services, or Oracle Messaging Gateway. You must also configure Oracle pre-compilers for your operating system and if desired, configure Oracle Advanced Security.

The Companion CD contains additional Oracle Database 10g software that you may want to install to improve performance, for example, for using Oracle JVM, Oracle *interMedia* or Oracle Text.

The Oracle Universal Installer and Real Application Clusters

The Oracle Universal Installer (OUI) is a graphical user interface that facilitates the installation of Cluster Ready Services (CRS) and Oracle Database 10g software. When the OUI installs the Oracle software, Oracle recommends that you select a preconfigured database or use the Database Configuration Assistant (DBCA) interactively to create your cluster database. You can also manually create your database as described in procedures posted at <http://otn.oracle.com>. Oracle recommends that you use Automatic Storage Management (ASM). If you are not using ASM or if you are not using a cluster file system, then configure shared raw devices before you create your database.

When you install CRS or RAC, the OUI copies the Oracle software onto the node from which you are running it. If your Oracle home is not on a cluster file system, then the OUI propagates the software onto the other nodes that you have selected to be part of your OUI installation session.

When you are installing the database software, the OUI also installs the software for Oracle Enterprise Manager Database Control and integrates this tool into your cluster environment. Once installed, Enterprise Manager Database Control is fully configured and operational for RAC. You can also install Enterprise Manager Grid Control onto other client machines outside your cluster to monitor multiple RAC and single-instance Oracle database environments.

The Enterprise Manager Agent is installed when you install Enterprise Manager. The Agent must be running on each node that is part of your RAC database so that Enterprise Manager can automatically discover and manage your RAC database and its instances.

See Also:

- The *Oracle Universal Installer Concepts Guide* for more details about the OUI
- The *Oracle Real Application Clusters Administrator's Guide* for information about using Enterprise Manager to administer RAC environments

Storage Considerations for Installing Oracle Database 10g Real Application Clusters

Oracle recommends using Automatic Storage Management (ASM) or a cluster file system with Oracle Managed Files (OMF) for database storage. This section provides an overview of ASM.

Overview of Automatic Storage Management

You can use ASM to simplify the administration of Oracle database files. Instead of having to manage potentially thousands of database files, using ASM, you need to manage only a small number of disk groups. A disk group is a set of disk devices that ASM manages as a single logical unit. You can define a particular disk group as the default disk group for a database, and Oracle will automatically allocate storage for, create, or delete, the files associated with the appropriate database object. When administering the database, you need only refer to database objects by name, rather than by file name.

When using ASM with a single Oracle home for database instances on a node, the ASM instance can run from that same home. If you are using ASM with Oracle database instances from multiple database homes on the same node, then Oracle recommends that you run the ASM instance from an Oracle home that is distinct from the database homes. In addition, the ASM home should be installed on every cluster node. This prevents the accidental removal of ASM instances that are in use by databases from other homes during the de-installation of a database's Oracle home.

Benefits of Oracle Storage Management

ASM provides many of the same benefits as storage technologies such as RAID or logical volume managers (LVMs). Like these technologies, ASM enables you to create a single disk group from a collection of individual disk devices. It balances I/O to the disk group across all of the devices in the disk group. It also implements striping and mirroring to improve I/O performance and data reliability.

However, unlike RAID or LVMs, ASM implements striping and mirroring at the file level. This implementation enables you to specify different storage attributes for individual files in the same disk group.

Disk Groups and Failure Groups

A disk group can include any number of disk devices. Each disk device can be an individual physical disk, a multiple disk device such as a RAID storage array or logical volume, or even a partition on a physical disk. However, in most cases, disk groups consist of one or more individual physical disks. To enable ASM to balance I/O and storage appropriately within the disk group, all devices in the disk group should have similar, if not identical, storage capacity and performance.

Note: Do not specify more than one partition on a single physical disk as a disk group device. 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 ASM disk group, Oracle does not recommend their use. Because logical volume managers can hide the physical disk architecture, ASM may not operate effectively when logical volumes are specified as disk group devices.

When you add a device to a disk group, you can specify a failure group for that device. Failure groups identify disk devices that have common failure characteristics, for example, devices that are attached to the same controller. If the controller fails, then all devices attached to it become unavailable. By default, each device also belongs to its own failure group. By using the failure groups you specify, ASM can distribute data among the devices in the disk group to help minimize the risk of data loss caused by component failures.

Redundancy Levels

ASM provides three levels of mirroring, called redundancy levels, that you can specify when creating a disk group. The redundancy levels are:

- External redundancy

In disk groups created with external redundancy, the contents of the disk group are not mirrored by ASM. You might choose this redundancy level when:

- The disk group contains devices, such as RAID devices, that provide their own data protection
- Your use of the database does not require uninterrupted access to data, for example, in a development environment where you have a suitable back-up strategy

- Normal redundancy

In disk groups created with normal redundancy, the contents of the disk group are two-way mirrored by default. However, you can choose to create certain files that are not mirrored. To create a disk group with normal redundancy, you must specify at least two failure groups (a minimum of two devices).

The effective disk space of a disk group that uses normal redundancy is half the total disk space of all of its devices.

- High redundancy

In disk groups created with high redundancy, the contents of the disk group are three-way mirrored by default. However, you can choose to create certain files that are only two-way mirrored, or that are not mirrored. To create a disk group with high redundancy, you must specify at least three failure groups (a minimum of three devices).

The effective disk space of a disk group that uses high redundancy is one-third of the total disk space of all of its devices.

ASM and Installation Types

The type and number of disk groups that you can create when installing Oracle software depends on the type of database you choose to create during the installation, as follows:

- Preconfigured database

If you choose to create the default preconfigured database that uses ASM, the OUI prompts you to specify two disk device names, which it uses to create a disk group named DATA, with normal redundancy.
- Custom database

If you choose to create a custom database that uses ASM, you can create one or more disk groups. These disk groups can use one or more devices. For each disk group, you can specify the redundancy level that suits your requirements.

The following table lists the total disk space required in all disk group devices for a typical preconfigured database, depending on the redundancy level you choose to use for the disk group:

Redundancy Level	Total Disk Space Required
External	1 GB
Normal	2 GB (on a minimum of two devices)
High	3 GB (on a minimum of three devices)

Additional Considerations for Using Oracle Database 10g Features in RAC

Oracle recommends using the following Oracle Database 10g features to simplify RAC database management:

- Enterprise Manager—Use Enterprise Manager to administer your *entire* processing environment, not just the RAC database. Enterprise Manager enables you to manage a RAC database with its instance targets, listener targets, host targets, and a cluster target, as well as ASM targets if you are using ASM storage for your database.
- Automatic undo management—Automatically manages undo processing.
- Automatic segment-space management—Automatically manages segment freelists and freelist groups.
- Locally managed tablespaces—Enhances space management performance.

See Also: The *Oracle Real Application Clusters Administrator's Guide* for more information about these features in RAC environments

Oracle Database 10g and Real Application Clusters

Oracle Database 10g provides single-instance database software and the additional components to operate RAC databases. Some of the RAC-specific components include:

- Cluster Ready Services (CRS)
- A RAC-enabled Oracle home

The Cluster Ready Services Clusterware

The OUI installs CRS on each node on which the OUI detects that vendor clusterware is running. If vendor clusterware is not running, then you must use the OUI to select the nodes on which to install CRS. The CRS home can be either shared by all nodes or private to each node depending on your responses when you run the OUI. The home that you select for CRS *must* be different from the RAC-enabled Oracle home.

When vendor clusterware is present, CRS interacts with the vendor clusterware to coordinate cluster membership information. For Oracle Database 10g on Linux, CRS coexists with but does not interact with previous Oracle clusterware versions. In addition:

- You may use vendor clusterware for all operating systems except Linux.
- For SunClusters, Oracle provides a UDLM patch that you must install onto each node in the cluster from the `/racpatch` directory on the Oracle Cluster Ready Services Release 1 (10.1.0.2) CD-ROM before installing and configuring RAC. Although you may have an operative pre-Oracle 10g version of the UDLM, you must install the Oracle 10g UDLM.

The Installed Real Application Clusters Components

All instances in RAC environments share the control file, server parameter file, redo log files, and all datafiles. These files reside on a shared cluster file system or on shared disks. Either of these types of file configurations are accessed by all the cluster database instances. Each instance also has its own set of redo log files. During failures, shared access to redo log files enables surviving instances to perform recovery.

Oracle Database 10g Real Application Clusters Version Compatibility

You can install and operate different versions of Oracle cluster database software on the same computer as described in the following points:

- If you install RAC on a system that already has an Oracle Database 10g RAC installation, then the OUI prompts you to install additional Oracle Database 10g products if you have not already installed all of them.
- You can also install multiple Oracle Database 10g RAC homes on the same node. The multiple homes feature enables you to install one or more releases on the same machine in multiple Oracle home directories. However, each node can have only one CRS home.
- The OUI also enables you to uninstall and re-install Oracle Database 10g Real Application Clusters if needed.
- In addition, you cannot install Oracle Database 10g RAC into an existing Oracle home. If you have an Oracle home for Oracle Database 10g or an Oracle home for an earlier Oracle cluster database software release, then you must use a different home for each subsequent installation. If the OUI detects an earlier version, then it runs the DBUA to prompt you for upgrade information. In this case, select the **Upgrade** option to upgrade to Oracle Database 10g release 1 (10.1).

Note: Do not move Oracle binaries from one Oracle home to another because this causes dynamic link failures.

Note: On some platforms, pre-Oracle Database 10g cluster manager implementations were referred to as "Cluster Manager". In Oracle Database 10g, the cluster manager role is performed by Cluster Synchronization Services (CSS) on all platforms. The OCSSD performs this function.

Required Groups

Depending on whether this is the first time you are installing the Oracle server software on your system, you may need to create several groups and a user as described later in the pre-installation procedures. The required groups and user are:

- The Oracle Inventory group (`oinstall`)

You must create this group the first time you install Oracle software on the system. The usual name for this group is `oinstall`. This group owns the Oracle inventory which is a catalog of all of the Oracle software installed on the system.
- The OSDBA group (`dba`)

You must create this group the first time you install Oracle software on the system. It identifies users that have database administrative privileges (the `SYSDBA` and `SYSOPER` privileges). The default name for this group is `dba`. To specify a group name other than the default, choose the **Custom** installation type to install the software. You must also create an OSDBA group if an OSBDA group exists but you want to give a different group of users database administrative privileges in a new Oracle server installation.
- The OSOPER group (`oper`)

This is an optional group. Create this group if you want a separate group of users to have a limited set of database administrative privileges (the `SYSDBA` privilege). The default name for this group is `oper`. To use this group, choose the **Custom** installation type to install the software. To use an OSOPER group, you must create it in the following circumstances:

 - If an OSOPER group does not exist, for example, if this is the first installation of Oracle server software on the system
 - If an OSOPER group exists but you want to give a different group of users database operator privileges in a new Oracle server installation
- The Oracle Software Owner user (`oracle`)

You must create this user the first time you install Oracle software on the system. This user owns all of the software installed during the installation. The usual name for this user is `oracle`. This user must have the Oracle Inventory group as its primary group and the OSDBA group as its secondary group. It must also have the OSOPER group as a secondary group if you choose to create that group. If an Oracle software owner user exists, but you want to use a different user, with different group membership, in a new Oracle server installation, you must give database administrative privileges to those groups.

A single Oracle Inventory group is required for all installations of Oracle software on the system. However, you can create different Oracle software owner users, OSDBA groups, and OSOPER groups (other than `oracle`, `dba`, and `oper`) for separate installations. Using different groups enables you to grant DBA privileges to a particular operating system user on one database, which they would not have on another database on the same system.

See Also: The *Oracle Database Administrator's Reference, 10g Release 1 (10.1) for UNIX Systems* and the *Oracle Database 10g Administrator's Guide* contain more information about the OSDBA and OSOPER groups and the SYSDBA and SYSOPER privileges.

Part II

Real Application Clusters Platform-Specific Pre-Installation Procedures

Part II describes the platform-specific pre-installation procedures for Oracle Database 10g Real Application Clusters (RAC). The chapters in Part II are:

- [Chapter 2, "Pre-Installation Tasks for Installing RAC on AIX-Based Systems"](#)
- [Chapter 3, "Pre-Installation Tasks for Installing RAC on HP-UX Systems"](#)
- [Chapter 4, "Pre-Installation Tasks for Installing RAC on HP Tru64 UNIX"](#)
- [Chapter 5, "Pre-Installation Tasks for Installing RAC on Linux-Based Systems"](#)
- [Chapter 6, "Pre-Installation Tasks for Installing RAC on Solaris Operating System \(SPARC 64-Bit\)"](#)

Pre-Installation Tasks for Installing RAC on AIX-Based Systems

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer. It includes information about the following tasks:

- Log In to the System as root
- Check the Hardware Requirements
- Check the Network Requirements
- Checking the Software Requirements
- Create Required UNIX Groups and User
- Configure Shell Limits and Kernel Parameters
- Identify Required Software Directories
- Identify or Create an Oracle Base Directory
- Configure Disk Storage for Oracle CRS and Database Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Configure Raw Logical Volumes
- Verifying the Cluster Software Configuration
- Stop Existing Oracle Processes
- Configure the oracle User's Environment

Log In to the System as root

Before you install the Oracle software, you must complete several tasks as the `root` user. To log in as the `root` user, complete one of the following procedures:

Note: Unless you intend to complete a silent installation, you must install the software from an X Window System workstation, an X terminal, or a PC or other system with X server software installed.

For more information about non-interactive and silent installations, see the appendices in *Oracle Database Installation Guide for UNIX Systems*.

- If you are installing the software from an X Window System workstation or X terminal:
 1. Start a local terminal session, for example, an X terminal (`xterm`).
 2. If you are not installing the software on the local system, enter the following command to enable remote hosts to display X applications on the local X server:

```
$ xhost +
```

3. If you want to install the software on a remote system, enter a command similar to the following to connect to that system:

```
$ telnet remote_host
```

4. If you are not logged in as the `root` user, enter the following command to switch user to `root`:

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

- If you are installing the software from a PC or other system with X server software installed:

Note: If necessary, see your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

1. Start the X server software.
2. Configure the security settings of the X server software to permit remote hosts to display X applications on the local system.
3. Connect to the remote system where you want to install the software and start a terminal session on that system, for example, an X terminal (`xterm`).
4. If you are not logged in as the `root` user on the remote system, enter the following command to switch user to `root`:

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

Check the Hardware Requirements

The system must meet the following minimum hardware requirements:

- 512 MB of physical RAM
- 1 GB of swap space (or twice the size of RAM)

On systems with 2 GB or more of RAM, the swap space can be between one and two times the size of RAM, up to a maximum of 2 GB of swap space. See the *Oracle Database Administrator's Reference for UNIX* for more information about configuring swap space.

- 400 MB of disk space in the `/tmp` directory
- Up to 4 GB of disk space for the Oracle software, depending on the installation type
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirements for databases that use Automatic Storage Management (ASM) or raw device storage is described later in this chapter.

Additional disk space, either on a file system or in an ASM disk group, is required for the flash recovery area if you choose to configure automated backups.

To ensure that the system meets these requirements, follow these steps:

1. To determine the physical RAM size, enter the following command:

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

If the size of the physical RAM installed in the system is less than the required size, you must install more memory before continuing.

2. To determine the size of the configured swap space, enter the following command:

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

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

3. To determine the amount of disk space available in the `/tmp` directory, enter the following command:

```
# df -k /tmp
```

If there is less than 400 MB of disk space available in the `/tmp` directory, complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to achieve the required disk space.
- Set the `TEMP` and `TMPDIR` environment variables when setting the `oracle` user's environment (described later).
- Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.

4. To determine the amount of free disk space on the system, enter the following command:

```
# df -k
```

5. To determine whether the system architecture can run the software, enter the command:

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

The expected output is the value 64.

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

Check the Network Requirements

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Real Application Clusters installation.

Note: For the most up-to-date information about supported network protocols and hardware for RAC installations, see the Certify pages on the Oracle*MetaLink* Web site:

<http://metalink.oracle.com>

Network Hardware Requirements

Each node in the cluster must meet the following requirements:

- Each node must have at least two network adapters; one for the public network interface and one for the private network interface (the interconnect).
- The interface names associated with the network adapters for each network must be the same on all nodes.
- For increased reliability, you can configure redundant public and private network adapters for each node.
- For the public network, each network adapter must support TCP/IP.
- For the private network, the interconnect must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended).

Note: UDP is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS.

Token-Ring is not supported for the interconnect.

IP Address Requirements

Before starting the installation, you must identify or obtain the following IP addresses for each node:

- An IP address and an associated host name registered in the domain name service (DNS) for each public network interface

- One unused virtual IP address and an associated virtual host name registered in DNS that you will configure for the primary public network interface
The virtual IP address must be in the same subnet as the associated public interface. After installation, you can configure clients to use the virtual host name or IP address. If a node fails, its virtual IP address fails over to another node.
- A private IP address and optional host name for each private interface
Oracle recommends that you use private network IP addresses for these interfaces, for example: 10.*.* or 192.168.*.*. You can use the `/etc/hosts` file on each node to associate private host names with private IP addresses.

For example, if each node has two public and two private interfaces, you might have the following host names and IP addresses on one of the nodes (`rac1`) and similar host names and IP addresses on the other nodes:

Host Name	Type	IP Address	Registered In
<code>rac1.mydomain.com</code>	Public	143.47.43.100	DNS
<code>rac1-2.mydomain.com</code>	Public	143.46.51.101	DNS
<code>rac1-vip.mydomain.com</code>	Virtual	143.46.43.104	DNS
<code>rac1-priv1</code>	Private	10.0.0.1	<code>/etc/hosts</code>
<code>rac1-priv2</code>	Private	10.0.0.2	<code>/etc/hosts</code>

Checking the Network Requirements

To verify that each node meets the requirements, follow these steps:

1. If necessary, install the network adapters for the public and private networks and configure them with either public or private IP addresses.
2. Register the host names and IP addresses for the public network interfaces in DNS.
3. For each node, register one virtual host name and IP address in DNS.
4. For each private interface on every node, add a line similar to the following to the `/etc/hosts` file on all nodes, specifying the private IP address and associated private host name:

```
10.0.0.1    rac1-priv1
```

5. Verify that the entry for the loopback address 127.0.0.1 in the `/etc/hosts` file is similar to the following:

```
127.0.0.1  localhost  loopback
```

The name `localhost` must be the first host name specified on this line.

6. To identify the interface name and associated IP address for every network adapter, enter the following command:

```
# /etc/ifconfig -a
```

From the output, identify the interface name and IP address for all network adapters that you want to specify as public or private network interfaces.

Note: When you install Oracle CRS and RAC, you will require this information.

Checking the Software Requirements

See the following sections for information about checking the software requirements.

Note: The Oracle Universal Installer performs checks on your system to verify that it meets the requirements listed for AIX. To ensure that these checks pass, verify the requirements before you start the Installer.

- ["Check for Required Software"](#) on page 2-6
- ["Check for Required Patches"](#) on page 2-7

Check for Required Software

Depending on the products that you intend to install, verify that the following software is installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	<p>Operating system version and maintenance level: AIX 5L version 5.2, Maintenance Level 1 or higher</p> <hr/> <p>Operating system filesets: bos.adt.base bos.adt.lib bos.adt.libm bos.perf.libperfstat bos.perf.perfstat bos.perf.proctools</p>
Real Application Clusters	<p>High Availability Cluster Multi-Processing (HACMP) v5.1 Note: HACMP is required only if you want to use raw logical volumes for Oracle CRS or database files.</p> <p>General Parallel File System (GPFS) v2.1 Note: GPFS is required only if you want to use a cluster file system for Oracle CRS or database files.</p>
Oracle Spatial	<p>AIXwindows Motif Libraries: x11.motif.lib</p> <p>Note: This software is required only to build the sample programs.</p>
Oracle Messaging Gateways	<p>IBM WebSphere MQ V5.3, client and server: mqm.Client.Bnd mqm.Server.Bnd</p>

Installation Type or Product	Requirement
Oracle JDBC/OCI Drivers	<p>You can use the following optional IBM JDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation:</p> <ul style="list-style-type: none"> ■ JDK 1.4.1.1 (64-bit) ■ JDK 1.3.1.11 (32-bit) ■ JDK 1.2.2.18 <p>Note: IBM JDK 1.4.1.2 (32-bit) is installed with this release.</p>
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>C for AIX Compiler V6.0.0.4: vac.c</p> <p>VisualAge C++ Compiler V6.0: vacpp.cmp.core</p>
Oracle Transparent Gateway for Sybase	<p>Either of the following:</p> <ul style="list-style-type: none"> ■ Sybase Adaptive Server Enterprise V12 or 12.5 ■ Sybase Open Client V12 or 12.5 <p>Open Client is required if the Sybase server is not on the local node.</p>

To ensure that the system meets these requirements, follow these steps:

1. To determine which version of AIX is installed, enter the following command:

```
# oslevel -r
```

If the operating system version is lower than AIX 5.2.0.0 Maintenance Level 1 (5200-01), upgrade your operating system to this level. AIX 5L version 5.2 maintenance packages are available from the following Web site:

<https://techsupport.services.ibm.com/server/aix.fdc>

2. To determine whether the required file sets are installed and committed, enter a command similar to the following:

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

If a file set is not installed and committed, then install it. See your operating system or software documentation for information about installing file sets.

Check for Required Patches

Depending on the products that you intend to install, verify that the following patches are installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	<p>Authorized Problem Analysis Reports (APARs):</p> <ul style="list-style-type: none"> ■ IY43980: libperfstat.h not ANSI-compliant ■ IY44810: DSI IN BMRECYCLE ■ IY45462: Definition of isnan() in math.h incorrect ■ IY45707: J2 READAAHEAD/CIO INTERACTION ■ IY46214: dropping partial connections leaves them on so_q0 ■ IY46605: exec of 32 bit application can fail on 64 bit kernel ■ IY48525: SDK 1.4.1 32-BIT SR1: CA141-20030930 ■ IY51801: race condition in aio_nwait_timeout
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Program Technical Fixes (PTFs):</p> <ul style="list-style-type: none"> ■ U489726 - vac.C.6.0.0.4 (or later)
Oracle JDBC/OCI Drivers	<p>Note: These APARs are required only if you are using the associated JDK version.</p> <p>APARs required for JDK 1.4.1.1 (64-bit):</p> <ul style="list-style-type: none"> ■ IY48526: SDK 1.4.1 64-BIT SR1: CAIX64141-20030930 <p>APARs required for JDK 1.3.1.11 (32-bit):</p> <ul style="list-style-type: none"> ■ IY47055: SDK 1.3.1 32-BIT PTF: CA131-20030630A <p>APARs required for JDK 1.2.2.18:</p> <ul style="list-style-type: none"> ■ IY40034: SDK 1.2.2 PTF: CA122-20030115
Oracle Messaging Gateway	<p>Corrective service diskettes (CSDs) for WebSphere MQ: CSD03 or later for WebSphere MQ V5.3 (U485561)</p>
Real Application Clusters	<p>APARs and PTFs required for HACMP:</p> <ul style="list-style-type: none"> ■ IY42783: CT:LX: RMC daemon may hang if managed nodes recycle ■ IY43602: DISK FAILURES CAUSING QUORUM TO BE LOST IS NOT ■ IY45695: BASE FIXES FOR HACMP 5.1.0 ■ U496124: cluster.es.server.rte.5.1.0.2 <p>APARs and PTFs required for GPFS:</p> <ul style="list-style-type: none"> ■ IY52454 efix ■ U489058: mmfs.base.cmds.3.5.0.6 ■ U496347: mmfs.gpfs.rte.3.5.0.10 ■ U496395: mmfs.gpfs.rte.2.1.0.10

To ensure that the system meets these requirements, follow these steps:

1. To determine whether an APAR is installed, enter a command similar to the following:

```
# /usr/sbin/instfix -i -k "IY22854 IY26778 ..."
```

If an APAR is not installed, download it from the following Web site and install it:

<https://techsupport.services.ibm.com/server/aix.fdc>

If you are using GPFS and require efix IY52454, see the following FTP site for download instructions:

```
ftp://service.software.ibm.com/aix/efixes/iy52454/
```

2. To determine whether a PTF is installed, enter a command similar to the following:

```
# lspp -l -B U489726 U485561 ...
```

If a PTF is not installed, download it from the following Web site and install it:

```
https://techsupport.services.ibm.com/server/aix.fdc
```

3. If you require a CSD for WebSphere MQ, see the following Web site for download and installation information:

```
http://www.ibm.com/software/integration/mqfamily/support/summary/aix.html
```

4. To continue completing pre-installation tasks, go to the "[Create Required UNIX Groups and User](#)" section on page 2-9.

Create Required UNIX Groups and User

Depending on whether this is the first time Oracle software is being installed on this system and on the products that you are installing, you may need to create several UNIX groups and a UNIX user.

The following UNIX groups and user are required if you are installing Oracle Database:

- The OSDBA group (`dba`)

You must create this group the first time you install Oracle Database software on the system. It identifies UNIX users that have database administrative privileges (the SYSDBA privilege). The default name for this group is `dba`.

If you want to specify a group name other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of this group. In this case, the Installer prompts you to specify the name of this group.

- The OSOPER group (`oper`)

This is an optional group. Create this group if you want a separate group of UNIX users to have a limited set of database administrative privileges (the SYSOPER privilege). By default, members of the OSDBA group also have the SYSOPER privilege.

If you want to specify a separate OSOPER group, other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of the `dba` group. In this case, the Installer prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

- An unprivileged user (`nobody`)

You must verify that the unprivileged user `nobody` exists on the system. The `nobody` user must own the external jobs (`ext job`) executable after the installation.

The following UNIX group and user are required for all installations:

- The Oracle Inventory group (`oinstall`)

You must create this group the first time you install Oracle software on the system. The usual name chosen for this group is `oinstall`. This group owns the Oracle inventory, which is a catalog of all Oracle software installed on the system.

Note: If Oracle software is already installed on the system, the existing Oracle Inventory group must be the primary group of the UNIX user that you use to install new Oracle software. The following sections describe how to identify an existing Oracle Inventory group.

- The Oracle software owner user (`oracle`)

You must create this user the first time you install Oracle software on the system. This user owns all of the software installed during the installation. The usual name chosen for this user is `oracle`. This user must have the Oracle Inventory group as its primary group. It must also have the OSDBA and OSOPER groups as a secondary groups.

A single Oracle Inventory group is required for all installations of Oracle software on the system. After the first installation of Oracle software, you must use the same Oracle Inventory group for all subsequent Oracle software installations on that system. However, you can choose to create different Oracle software owner users, OSDBA groups, and OSOPER groups (other than `oracle`, `dba`, and `oper`) for separate installations. By using different groups for different installations, members of these different groups have DBA privileges only on the associated databases rather than on all databases on the system.

See Also: *The Oracle Database Administrator's Reference for UNIX Systems* and the *Oracle Database Administrator's Guide* contain more information about the OSDBA and OSOPER groups and the SYSDBA and SYSOPER privileges.

The following sections describe how to create the required UNIX user and groups.

Note: The following sections describe how to create local users and groups. As an alternative to creating local users and groups, you could create the appropriate users and groups in a directory service, for example, Network Information Services (NIS). For information about using directory services, contact your system administrator or see your operating system documentation.

Creating the Oracle Inventory Group

You must create the Oracle Inventory group if it does not already exist. The following subsections describe how to determine the Oracle Inventory group name, if it exists, and how to create it if necessary.

Determine Whether the Oracle Inventory Group Exists

When you install Oracle software on the system for the first time, the Installer creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group and the path of the Oracle Inventory directory. To determine whether the Oracle Inventory group exists, enter the following command:

```
# more /etc/oraInst.loc
```

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inst_group` parameter shows the name of the Oracle Inventory group (`oinstall`).

Create the Oracle Inventory Group

If the `oraInst.loc` file does not exist, create the Oracle Inventory group as follows:

1. Enter the following command:

```
# smit security
```

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

Creating the OSDBA Group

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system
- An OSBDA group exists, but you want to give a different group of UNIX users database administrative privileges in a new Oracle installation

If the OSDBA group does not exist or if you require a new OSDBA group, create it as follows, using the group name `dba` unless a group with that name already exist:

1. Enter the following command:

```
# smit security
```

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

Creating an OSOPER Group (Optional)

Create an OSOPER group only if you want to identify a group of UNIX users with a limited set of database administrative privileges (SYSOPER operator privileges). For most installations, it is sufficient to create only the OSDBA group. If you want to use an OSOPER group, you must create it in the following circumstances:

- If an OSOPER group does not exist, for example, if this is the first installation of Oracle Database software on the system
- If an OSOPER group exists, but you want to give a different group of UNIX users database operator privileges in a new Oracle installation

If you require a new OSOPER group, create it as follows, using the group name `oper` unless a group with that name already exists:

1. Enter the following command:

```
# smit security
```

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

Creating the Oracle Software Owner User

You must create an Oracle software owner user in the following circumstances:

- If an Oracle software owner user does not exist, for example, if this is the first installation of Oracle software on the system
- If an Oracle software owner user exists, but you want to use a different UNIX user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation

Determine Whether an Existing Oracle Software Owner User Exists

To determine whether an Oracle software owner user named `oracle` exists, enter the following command:

```
# id oracle
```

If the `oracle` user exists, the output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

If the user exists, determine whether you want to use the existing user or create a new user. If you want to use the existing user, ensure that the user's primary group is the Oracle Inventory group and that it is a member of the appropriate OSDBA and OSOPER groups. See one of the following sections for more information:

Note: If necessary, contact your system administrator before using or modifying an existing user.

- If you want to use the existing Oracle software owner user, and the user's primary group is the Oracle Inventory group, see ["Verifying That the UNIX User nobody Exists"](#) section on page 2-13.
- To modify an existing user, see ["Modifying an Existing Oracle Software Owner User"](#) section on page 2-13.
- To create a new user, see the following section.

Create a New Oracle Software Owner User

If the Oracle software owner user does not exist or if you require a new Oracle software owner user, create it as follows, using the user name `oracle` unless a user with that name already exists:

1. Enter the following command:

```
# smit security
```

2. Choose the appropriate menu items to create the `oracle` user, specifying the following information:
 - In the **Primary GROUP** field, specify the Oracle Inventory group, for example `oinstall`.
 - In the **Group SET** field, specify the OSDBA group and if required, the OSOPER group, for example `dba` or `dba, oper`.

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

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

```
# passwd oracle
```

See the "[Verifying That the UNIX User nobody Exists](#)" section on page 2-13 to continue.

Modifying an Existing Oracle Software Owner User

If the `oracle` user exists, but its primary group is not `oinstall` or it is not a member of the appropriate OSDBA or OSOPER groups, you can modify it as follows:

1. Enter the following command:


```
# smit security
```
2. Choose the appropriate menu items to modify the `oracle` user.
3. In the **Primary GROUP** field, specify the Oracle Inventory group, for example `oinstall`.
4. In the **Group SET** field, specify the secondary groups, for example `dba` and `oper`.
5. Press F10 to exit.

Verifying That the UNIX User nobody Exists

Before installing the software, verify that the UNIX user `nobody` exists on the system:

1. To determine whether the user exists, enter the following command:

```
# id nobody
```

If this command displays information about the `nobody` user, you do not have to create that user.

2. If the `nobody` user does not exist, enter the following command to create it:

```
# smit security
```

Specify the appropriate options to create an unprivileged `nobody` user, then press F10 to exit.

3. Repeat this procedure on all of the other cluster nodes.

Create Identical Users and Groups on Other Cluster Nodes

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

The Oracle software owner user and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes. To create these identical users and groups, you must identify the user ID and group IDs assigned them on the node

where you created them, then create the user and groups with the same name and ID on the other cluster nodes.

Identifying the User and Group IDs

To determine the user ID (UID) of the Oracle software owner user and the group IDs (GID) of the Oracle Inventory, OSDBA, and OSOPER groups, follow these steps:

1. Enter the following command:

```
# id oracle
```

The output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

2. From the output, identify the UID for the `oracle` user and the GIDs for the groups to which it belongs.

Create the User and Groups on the Other Cluster Nodes

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

1. Log in to the next cluster node as `root`.
2. Enter the following command:

```
# smit security
```
3. Choose the appropriate menu items to create or modify the `oinstall` and `dba` groups, and if required, the `oper` group.

In the **Group NAME** and **Group ID** fields, specify the group names and GIDs that you identified previously.

Note: If you cannot use the same group ID for a particular group on this node, identify a group ID that *is* available on all nodes and specify that ID for the group on all of the nodes.

4. Choose the appropriate menu items to create or modify the `oracle` user, specifying the following information:
 - In the **User ID** field, specify the user ID that you identified in the previous subsection.
 - In the **Primary GROUP** field, specify the Oracle Inventory group, for example `oinstall`.
 - In the **Group SET** field, specify the OSDBA group and if necessary, the OSOPER group, for example `dba` or `dba, oper`.

Note: If you cannot use the same user ID for the `oracle` user on this node, identify a user ID that *is* available on all nodes and specify that ID for the user on all of the nodes.

5. Press F10 to exit.
6. Set the password of the `oracle` user:

```
# passwd oracle
```

Set Up User Equivalence for rsh and rcp on All Cluster Nodes

Note: This section describes how to set up user equivalence for `rcp`, which the Installer uses when copying Oracle software to the other cluster nodes. If you prefer, you can configure the Secure Shell (SSH) tool suite, so that the Installer uses `scp` instead of `rcp`. See the SSH documentation for information about setting up user equivalence for `scp`.

Before you install and use Oracle Real Application clusters, you must set up user equivalence on all cluster nodes, as follows:

1. On all cluster nodes, edit the `/etc/hosts.equiv` file in any text editor and add entries similar to the following for all nodes in the cluster:

```
node1-public oracle
node1-public.domain oracle
node2-public oracle
node2-public.domain oracle
.
.
.
```

In this example:

- `node1_public` is the host name associated with the primary public network interface
- `domain` is the domain name of the system, for example `mydomain.com`

Note: Alternatively, you can create a `.rhosts` file using the same format in the `oracle` user's home directory on all nodes in the cluster.

2. To test user equivalence on all cluster nodes, including the first node, enter commands similar to the following from the node where you will run the Installer:

```
# su oracle
$ rsh nodename1 date
$ rsh nodename2 date
.
.
.
$ exit
```

Each command should return only the date setting on the specified node. If a node prompts for a password or displays any other text, then user equivalence is not set up correctly on that node. In this case, verify that the `/etc/hosts.equiv` (or `~oracle/.rhosts`) file has the correct entries.

Configure Shell Limits and Kernel Parameters

On AIX systems, you do not need to configure kernel parameters. However, Oracle recommends that you set shell limits and system configuration parameters as described in this section. This section contains the covers the following topics:

- [Configure Shell Limits](#) on page 2-16
- [Configure System Configuration Parameters](#) on page 2-16

Notes: The kernel parameter and shell limit values shown in the following sections are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

You must set the kernel parameters and shell limits on all cluster nodes.

Configure Shell Limits

Verify that the shell limits shown in the following table are set to the values shown. The procedure following the table describes how to verify and set the values.

Shell Limit (As Shown in smit)	Recommended Value
Soft FILE size	-1 (Unlimited)
Soft CPU time	-1 (Unlimited)
	Note: This is the default value.
Soft DATA segment	-1 (Unlimited)
Soft STACK size	-1 (Unlimited)

To view the current value specified for these shell limits, and to change them if necessary, follow these steps:

1. Enter the following command:


```
# smit chuser
```
2. In the **User NAME** field, enter the user name of the Oracle software owner, for example `oracle`.
3. Scroll down the list and verify that the value shown for the soft limits listed in the previous table is -1.

If necessary, edit the existing value.
4. When you have finished making changes, press F10 to exit.

Configure System Configuration Parameters

Verify that the maximum number of processes allowed per user is set to 2048 or greater. The procedure following the table describes how to verify and set the value.

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

1. Enter the following command:

```
# smit chgsys
```
2. Verify that the value shown for **Maximum number of PROCESSES allowed per user** is greater than or equal to 2048.
 If necessary, edit the existing value.
3. When you have finished making changes, press F10 to exit.

Identify Required Software Directories

You must identify or create four directories for the Oracle software, as follows:

- Oracle base directory
- Oracle Inventory directory
- CRS home directory
- Oracle home directory

The following subsections describe the requirements for these directories.

Oracle Base Directory

The Oracle base directory acts as a top-level directory for Oracle software installations. It is analogous to the `C:\Oracle` directory used for Oracle software installations on Windows systems. On UNIX systems, the Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to the following for the Oracle base directory:

```
/mount_point/app/oracle_sw_owner
```

In this example:

- `mount_point` is the mount point directory for the file system that will contain the Oracle software.
 The examples in this guide use `/u01` for the mount point directory. However, you could choose another mount point directory, `/oracle` or `/opt/oracle` for example.
- `oracle_sw_owner` is the UNIX user name of the Oracle software owner, for example `oracle`.

You can use the same Oracle base directory for more than one installation or you can create separate Oracle base directories for different installations. If different UNIX users install Oracle software on the same system, each user must create a separate Oracle base directory. The following example Oracle base directories could all exist on the same system:

```
/u01/app/oracle
/u01/app/orauser
/opt/oracle/app/oracle
```

The following sections describe how to identify existing Oracle base directories that might be suitable for your installation and how to create a new Oracle base directory if necessary.

Regardless of whether you create a new Oracle base directory or decide to use an existing one, you must set the `ORACLE_BASE` environment variable to specify the full path to this directory.

Note: The Oracle base directory can be on a local file system, or on a supported cluster file system.

Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system. It is required by, and shared by, all Oracle software installations on a single system. The first time you install Oracle software on a system, the Installer prompts you to specify the path to this directory. If you are installing the software on a local file system, Oracle recommends that you choose the following path:

`oracle_base/oraInventory`

If the Oracle base directory is on a cluster file system, you must specify a path for the Oracle Inventory directory on a local file system, to enable all of the nodes to have separate inventories.

The Installer creates the directory that you specify and sets the correct owner, group, and permissions on it. You do not need to create it.

Note: All Oracle software installations rely on this directory. Make sure that you back it up regularly.

Do not delete this directory unless you have completely removed all Oracle software from the system.

CRS Home Directory

The CRS home directory is the directory where you choose to install the software for Oracle Cluster Ready Services. You must install CRS in a separate home directory. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the CRS home directory:

`oracle_base/product/10.1.0/crs`

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Oracle Home Directory

The Oracle home directory is the directory where you choose to install the software for a particular Oracle product. You must install different Oracle products, or different releases of the same Oracle product, in separate Oracle home directories. When you run the Installer, it prompts you to specify the path to this directory, as well as a name

that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the Oracle home directory:

```
oracle_base/product/10.1.0/db_1
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Identify or Create an Oracle Base Directory

Before starting the installation, you must either identify an existing Oracle base directory or if required, create a new one. This section contains information about the following:

- [Identifying an Existing Oracle Base Directory](#)
- [Creating a New Oracle Base Directory](#)

Note: You can choose to create a new Oracle base directory, even if other Oracle base directories exist on the system.

Identifying an Existing Oracle Base Directory

Existing Oracle base directories might not have paths that comply with OFA guidelines. However, if you identify an existing Oracle Inventory directory or existing Oracle home directories, you can usually identify the Oracle base directories, as follows:

- Identify an existing Oracle Inventory directory

Enter the following command on all cluster nodes to view the contents of the `oraInst.loc` file:

```
# more /etc/oraInst.loc
```

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inventory_loc` parameter identifies the Oracle Inventory directory (`oraInventory`) on that system. The parent directory of the `oraInventory` directory is typically an Oracle base directory. In the previous example, `/u01/app/oracle` is an Oracle base directory.

- Identify existing Oracle home directories

Enter the following command on all cluster nodes to view the contents of the `oratab` file:

```
# more /etc/oratab
```

If the `oratab` file exists, it contains lines similar to the following:

```
*:/u03/app/oracle/product/10.1.0/db_1:N
*/opt/orauser/infra_904:N
*/oracle/9.2.0:N
```

The directory paths specified on each line identify Oracle home directories. Directory paths that end with the user name of the Oracle software owner that you want to use are valid choices for an Oracle base directory. If you intend to use the `oracle` user to install the software, you could choose one of the following directories from the previous example:

```
/u03/app/oracle
/oracle
```

Note: If possible, choose a directory path similar to the first (`/u03/app/oracle`). This path complies with the OFA guidelines.

Before deciding to use an existing Oracle base directory for this installation, make sure that it satisfies the following conditions:

- It should not be on the same file system as the operating system.
- It must have an identical path on all cluster nodes or it must be on a supported cluster file system.
If you are not using a cluster file system, create identical Oracle base directories on the other nodes.
- It must have an identical path on all cluster nodes or it must be on a supported cluster file system.
If you are not using a cluster file system, create identical Oracle base directories on the other nodes.
- It must have sufficient free disk space on all cluster nodes, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB

Use the following command to determine the free disk space on the file system where the Oracle base directory is located:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the ["Creating a New Oracle Base Directory"](#) section on page 2-21.
- If an Oracle base directory exists and you want to use it, see the ["Configure Disk Storage for Oracle CRS and Database Files"](#) section on page 2-21.

When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify the directory you chose.

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the following section.

Creating a New Oracle Base Directory

Before you create a new Oracle base directory, you must identify an appropriate file system with sufficient free disk space, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB

To identify an appropriate file system, follow these steps:

1. Use the `df -k` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has appropriate free space.

Note: The file system can be a local file system or a supported cluster file system.

The path to the Oracle base directory must be the same on all nodes.

3. Note the name of the mount point directory for the file system that you identified.

To create the Oracle base directory and specify the correct owner, group, and permissions for it, follow these steps:

1. Enter commands similar to the following to create the recommended subdirectories in the mount point directory that you identified and set the appropriate owner, group, and permissions on them:

```
# mkdir -p /mount_point/app/oracle_sw_owner
# chown -R oracle:oinstall /mount_point/app
# chmod -R 775 /mount_point/app
```

If the mount point you identified is `/u01` and `oracle` is the user name of the Oracle software owner, the recommended Oracle base directory path is as follows:

```
/u01/app/oracle
```

2. If necessary, repeat the commands listed in the previous step to create the same directory on the other cluster nodes.
3. When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify this directory.

Configure Disk Storage for Oracle CRS and Database Files

This section describes the storage options available to you for storing Oracle Cluster Ready Services files, Oracle Database Files, and optionally, Oracle Database recovery files. After you choose the storage method that you want to use for each file type, see the following sections for information about configuring the required storage.

Note: You do not have to use the same storage option for each type of file.

Choosing a Storage Option for Oracle CRS Files

Before you install Oracle Cluster Ready Services, you must choose the storage option that you want to use for the Oracle Cluster Registry (100 MB) and CRS voting disk (20 MB). You cannot use Automatic Storage Management to store these files, because they must be accessible before any Oracle instance starts.

Apart from ASM, you can use any of the storage options listed in the next section for the Oracle CRS files.

Note: If you are not using HACMP, you must use a GPFS file system to store the Oracle CRS files.

Choosing a Storage Option for Oracle Database Files

If you want to create a database during the installation, you must choose one of the following storage options for the database files:

- Cluster file system using GPFS for AIX
- Automatic Storage Management
- Raw logical volumes using HACMP

Note: For the most up-to-date information about supported storage options for RAC installations, see the Certify pages on the *OracleMetaLink* Web site:

<http://metalink.oracle.com>

Choosing a Storage Option for Oracle Database Recovery Files

Additionally, if you want to enable automated backups during the installation, you must choose one of the following storage options for recovery files (the flash recovery area):

- Cluster file system using GPFS for AIX
- Automatic Storage Management

The storage option that you choose for recovery files can be the same as or different to the option you choose for the database files.

Configuring Disk Storage

For information about how to configure disk storage before you start the installation, see one of the following sections depending on your choice:

- To use a file system for Oracle CRS, database, or recovery file storage, see the "[Create Directories for Oracle CRS, Database, or Recovery Files](#)" section on page 2-23.
- To use ASM for database or recovery file storage, see the "[Configure Disks for Automatic Storage Management](#)" section on page 2-26.

- To use raw devices (partitions or logical volumes) for Oracle CRS or database file storage, see the "[Configure Raw Logical Volumes](#)" section on page 2-31.

Create Directories for Oracle CRS, Database, or Recovery Files

If you decide to place the Oracle CRS, database, or recovery files on a file system, use the following guidelines when deciding where to place them.

Guidelines for Placing Oracle CRS Files on a File System

The Installer does not suggest a default location for the Oracle Cluster Registry (OCR) or the Oracle CRS voting disk. If you choose to create these files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk.
- It must have at least 100 MB of free disk space for the OCR and 20 MB of free disk space for the CRS voting disk.
- For improved reliability, you should choose a file system on a highly available storage device, for example, a RAID device that implements mirroring.
- If you are placing the Oracle Cluster Ready Services software on a shared file system, you can use the same file system for these files.
- The `oracle` user must have write permissions to create the files in the path you specify.

Note: If you are upgrading from Oracle9i Release 2, you can continue to use the raw device or shared file that you used for the SRVM configuration repository instead of creating a new file for the OCR.

Guidelines for Placing Oracle Database Files on a File System

If you choose to place the Oracle database files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk.
- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.

This default location is not recommended for production databases.

- You can choose either a single file system or more than one file system to store the database files:
 - If you want to use a single file system, choose a file system on a physical device that is dedicated to the database.

For best performance and reliability, choose a RAID device or a logical volume on more than one physical device and implement the stripe-and-mirror-everything (SAME) methodology.
 - If you want to use more than one file system, choose file systems on separate physical devices that are dedicated to the database.

This method enables you to distribute physical I/O and create separate control files on different devices for increased reliability. It also enables you to fully implement the OFA guidelines. You must choose either the Advanced database creation option or the Custom installation type during the installation to implement this method.

- If you intend to create a preconfigured database during the installation, the file system (or file systems) that you choose must have at least 1.2 GB of free disk space.

For production databases, you must estimate the disk space requirement depending on the use you want to make of the database.

- For optimum performance, the file systems that you choose should be on physical devices that are used only by the database.
- The `oracle` user must have write permissions to create the files in the path you specify.

Guidelines for Placing Oracle Recovery Files on a File System

Note: You must choose a location for recovery files only if you intend to enable automated backups during the installation.

If you choose to place the Oracle recovery files on a file system, use the following guidelines when deciding where to place them:

- To prevent disk failure from making both the database files and the recovery files unavailable, place the recovery files in a file system on a different physical disk from the database files.

Note: Alternatively use an ASM disk group with a normal or high redundancy level for either or both file types.

- You must choose a shared file system, for example, a cluster file system on a shared disk.
- The file system that you choose should have at least 2 GB of free disk space.

The disk space requirement is the default disk quota configured for the flash recovery area (specified by the `DB_RECOVERY_FILE_DEST_SIZE` initialization parameter).

If you choose the Custom installation type or the Advanced database configuration option, you can specify a different disk quota value. After you create the database, you can also use Oracle Enterprise Manager Grid Control or Database Control to specify a different value.

For more information about sizing the flash recovery area, see the *Oracle Backup and Recovery Basics* manual.

- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.

This default location is not recommended for production databases.

- The `oracle` user must have write permissions to create the files in the path you specify.

Creating Required Directories

Note: You must complete this procedure only if you want to place the Oracle CRS, database, or recovery files on a separate file system to the Oracle base directory.

To create directories for the Oracle CRS, database, or recovery files on separate file systems to the Oracle base directory, follow these steps:

1. If necessary, configure the shared file systems that you want to use and mount them on each node.

Note: Make sure that the file systems are configured to mount automatically when a node reboots.

2. Use the `df -k` command to determine the free disk space on each mounted file system.
3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
CRS files	Choose a file system with at least 120 MB of free disk space
Datafiles	Choose either: <ul style="list-style-type: none"> ■ A single file system with at least 1.2 GB of free disk space ■ Two or more file systems with at least 1.2 GB of free disk space in total
Recovery files	Choose a file system with at least 2 GB of free disk space.

If you are using the same file system for more than one type of file, add the disk space requirements for each type to determine the total disk space requirement.

4. Note the names of the mount point directories for the file systems that you identified.
5. Enter commands similar to the following to create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:

- CRS file directory:

```
# mkdir /mount_point/oracrs
# chown oracle:oinstall /mount_point/oracrs
# chmod 775 /mount_point/oracrs
```

- Database file directory:

```
# mkdir /mount_point/oradata
# chown oracle:oinstall /mount_point/oradata
# chmod 775 /mount_point/oradata
```

- Recovery file directory (flash recovery area):

```
# mkdir /mount_point/flash_recovery_area
# chown oracle:oinstall /mount_point/flash_recovery_area
# chmod 775 /mount_point/flash_recovery_area
```

6. If you also want to use ASM or raw devices for storage, see one of the following sections:

- [Configure Disks for Automatic Storage Management](#)
- [Configure Raw Logical Volumes](#)

Otherwise see the "[Verifying the Cluster Software Configuration](#)" section on page 2-39.

Configure Disks for Automatic Storage Management

This section describes how to configure disks for use with ASM. Before you configure the disks, you must determine the number of disks and the amount of free disk space that you require. The following sections describe how to identify the requirements and configure the disks:

- [Identifying Storage Requirements for ASM](#)
- [Using an Existing ASM Disk Group](#)
- [Configuring Disks for ASM](#)

Note: Although this section refers to disks, you can also use zero-padded files on a certified NAS storage device in an ASM disk group. See the appendices in *Oracle Database Installation Guide for UNIX Systems* for information about creating and configuring NAS-based files for use in an ASM disk group.

Identifying Storage Requirements for ASM

To identify the storage requirements for using ASM, you must determine how many devices and the amount of free disk space that you require. To complete this task, follow these steps:

1. Determine whether you want to use ASM for Oracle database files (datafiles), recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. One can use the file system, while the other uses ASM.

If you choose to enable automated backups and you do not have a cluster file system available, you must choose ASM for recovery file storage.

If you enable automated backups during the installation, you can choose ASM as the storage mechanism for recovery files by specifying an ASM disk group for the flash recovery area. Depending how you choose to create a database during the installation, you have the following options:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can

decide whether you want to use the same ASM disk group for datafiles and recovery files, or you can choose to use different disk groups for each file type.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must use the same ASM disk group for database files and recovery files.
2. Choose the ASM redundancy level that you want to use for the ASM disk group.

The redundancy level that you choose for the ASM disk group determines how ASM mirrors files in the disk group and determines the number of disks and amount of disk space that you require, as follows:

- 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 ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use only RAID or similar devices that provide their own data protection mechanisms as disk devices in this type of disk group.

- Normal redundancy

In a normal redundancy disk group, ASM uses two-way mirroring by default, to increase performance and reliability. 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 most installations, Oracle recommends that you use normal redundancy disk groups.

- High redundancy

In a high redundancy disk group, 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.

While high redundancy disk groups do provide a high level of data protection, you must consider the higher cost of additional storage devices before deciding to use this redundancy level.

3. Determine the total amount of disk space that you require for the datafiles and recovery files.

Use the following table to determine the minimum number of disks and the minimum disk space requirements for the installation:

Redundancy Level	Minimum Number of Disks	Datafiles	Recovery Files	Both File Types
External	1	1.15 GB	2.3 GB	3.45 GB
Normal	2	2.3 GB	4.6 GB	6.9 GB
High	3	3.45 GB	6.9 GB	10.35 GB

You must also add additional disk space for the ASM metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

$$15 + (2 * \textit{number_of_disks}) + (126 * \textit{number_of_ASM_instances})$$

For example, for a four-node RAC installation, using three disks in a high redundancy disk group, you require an additional 525 MB of disk space:

$$(15 + (2 * 3) + (126 * 4)) = 525$$

If an existing ASM instance exists on the system, 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 installation.

The following section describes how to identify existing disk groups and determine the free disk space that they contain.

4. Optionally identify failure groups for the ASM disk group devices.

Note: You need to complete this step only if you intend to use an installation method that runs DBCA in interactive mode, for example, if you intend to choose the Custom installation type or the Advanced database configuration option. Other installation types do not enable you to specify failure groups.

If you intend to use a normal or high redundancy disk group, 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 SCSI controller, the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To avoid failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

5. If you are sure that a suitable disk group does not exist on the system, 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 ASM disk group should be the same size and have the same performance characteristics.
 - Do not specify more than one partition on a single physical disk as a disk group device. 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 ASM disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing ASM from optimizing I/O across the physical devices.

For information about completing this task, see the ["Configuring Disks for ASM"](#) section on page 2-30.

Using an Existing ASM Disk Group

If you want to use ASM as the storage option for either database or recovery files, and an existing ASM disk group exists, you have the following choices, depending on the installation method that you select:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to create a new disk group or use an existing one.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must choose an existing disk group for the new database; you cannot create a new disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The ASM instance that manages the existing disk group can be running in a different Oracle home directory.

To determine whether an existing ASM disk group exists, or to determine whether there is sufficient disk space in a disk group, you can use Oracle Enterprise Manager Database Grid Control or Database Control. Alternatively, you can use the following procedure:

1. View the contents of the `oratab` file to determine whether an ASM instance is configured on the system:

```
# more /etc/oratab
```

If an ASM instance is configured on the system, the `oratab` file should contain a line similar to the following:

```
+ASM:oracle_home_path:N
```

In this example, `+ASM` is the system identifier (SID) of the ASM instance and `oracle_home_path` is the Oracle home directory where it is installed. By convention, the SID for an ASM instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the ASM instance that you want to use.
3. Connect to the ASM instance as the `SYS` user with `SYSDBA` privilege and start the instance if necessary:

```
# $ORACLE_HOME/bin/sqlplus "SYS/SYS_password as SYSDBA"
SQL> STARTUP
```

4. Enter the following command to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

```
SQL> SELECT NAME,TYPE,TOTAL_MB,FREE_MB FROM V$ASM_DISKGROUP;
```

5. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.

6. If necessary, install or identify the additional disk devices required to meet the storage requirements listed in the previous section.

Note: If you are adding devices to an existing disk group, Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Configuring Disks for ASM

To configure disks for use with ASM, follow these steps:

1. If necessary, install the shared disks that you intend to use for the ASM disk group and reboot the system.

2. To make sure that the disks are available, enter the following command:

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

The output from this command is similar to the following:

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

3. If a disk is not listed as available, enter the following command to configure the new disks:

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

4. To determine which disks are configured in a volume group, enter the following command:

```
# /usr/sbin/lspv
```

The output from this command is similar to the following:

```
hdisk0      0000078752249812  rootvg
hdisk1      none              none
```

For each disk, this command shows:

- The disk device name
- Either the 16 character physical volume identifier (PVID) if the disk has one, or none
- Either the volume group to which the disk belongs, or none

The disks that you want to use might have a PVID, but they must not belong to a volume group.

5. If the disks that you want to use for the disk group do not have PVIDs, enter a command similar to the following for each disk that you want to use:

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

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

```
# chown oracle:dba /dev/rhdiskn
# chmod 660 /dev/rhdiskn
```

Note: If you are using a multi-pathing disk driver with ASM, ensure that you set the permissions only on the correct logical device name for the disk.

Configure Raw Logical Volumes

The following sections describe how to configure raw logical volumes:

- [Review Important Information](#)
- [Create a Volume Group](#)
- [Create Raw Logical Volumes in the New Volume Group](#)
- [Import the Volume Group on the Other Cluster Nodes](#)
- [Activate the Volume Group in Concurrent Mode on All Cluster Nodes](#) on page 2-37
- [Create the DBCA Raw Device Mapping File](#) on page 2-37

Note: If you are using ASM for database file storage, you need only create raw devices for the Oracle CRS files. However, Oracle recommends that you use a cluster file system to store the Oracle CRS files instead of using raw devices for them.

To use raw logical volumes for Oracle CRS and database storage, HACMP must be installed and configured on all cluster nodes.

Review Important Information

This section describes how to configure raw logical volumes for Oracle CRS and database file storage. The procedures in this section describe how to create a new volume group that contains the logical volumes required for both types of files.

Before you continue, review the following guidelines which contain important information about using volume groups with this release of Oracle Real Application Clusters:

- You must use concurrent-capable volume groups for Oracle CRS and database files.
- If you specify raw logical volumes for the Oracle Cluster Registry and Oracle CRS voting disk during the Oracle CRS installation, the installation process configures the Oracle CRS startup scripts to activate the volume group that contains these logical volumes when Oracle CRS starts.

Note: For this reason, you must not configure the volume group that contains the Oracle CRS files in a HACMP concurrent resource group.

- The Oracle CRS files require less than 200 MB of disk space. To make efficient use of the disk space in a volume group, Oracle recommends that you use the same volume group for the logical volumes for both the CRS files and the database files.

- If you are upgrading an existing Oracle9i release 2 RAC installation that uses raw logical volumes, you can use the existing SRVM configuration repository logical volume for the OCR and create a new logical volume in the same volume group for the Oracle CRS voting disk. However, you must remove this volume group from the HACMP concurrent resource group that activates it before you install Oracle CRS.

See Also: For information about removing a volume group from a concurrent resource group, see the HACMP documentation.

Note: If you are upgrading a database, you must also create a new logical volume for the SYSAUX tablespace. See the "[Create Raw Logical Volumes in the New Volume Group](#)" section on page 2-34 for more information about the requirements for the Oracle CRS voting disk and SYSAUX logical volumes.

- You must use a HACMP concurrent resource group to activate new or existing volume groups that contain only database files (not Oracle CRS files).

See Also: For information about adding a volume group to a new or existing concurrent resource group, see the HACMP documentation.

- All volume groups that you intend to use for Oracle CRS or database files must be activated in concurrent mode before you start the installation.
- The procedures in this section describe how to create basic volumes groups and volumes. If you want to configure more complex volumes, using mirroring for example, use this section in conjunction with the HACMP documentation.

Create a Volume Group

To create a volume group for the Oracle database files and the Oracle CRS files, follow these steps:

1. If necessary, install the shared disks that you intend to use.
2. To make sure that the disks are available, enter the following command:

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

The output from this command is similar to the following:

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

3. If a disk is not listed as available, enter the following command to configure it:

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

4. To determine whether the disks are configured in a volume group, enter the following command:

```
# /usr/sbin/lspv
```

The output from this command is similar to the following:

```
hdisk0      0000078752249812  rootvg
```

```
hdisk1    none          none
hdisk4    00034b6fd4ac1d71  ccvg1
```

For each disk, this command shows:

- The disk device name
- Either the 16 character physical volume identifier (PVID) if the disk has one, or none
- Either the volume group to which the disk belongs, or none

The disks that you want to use might have a PVID, but they must not belong to existing volume groups.

5. If the disks that you want to use for the volume group do not have PVIDs, enter a command similar to the following for each disk that you want to use:

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

6. To identify used device major numbers, enter the following command on each node of the cluster:

```
# ls -la /dev | more
```

This command displays information about all configured devices, similar to the following:

```
crw-rw----  1 root    system   45,  0 Jul 19 11:56 vg1
```

In this example, 45 is the major number of the vg1 volume group device.

7. Identify an appropriate major number that is unused on all nodes in the cluster.
8. To create a volume group, enter a command similar to the following, or use SMIT (`smit mkvg`):

```
# /usr/sbin/mkvg -y VGname -B -s PPSize -V majornum -n \  
-C PhysicalVolumes
```

The following table describes the options and variables used in this example. See the `mkvg` man page for more information about these options.

Command Option	SMIT Field	Sample Value and Description
-y <i>VGname</i>	VOLUME GROUP name	oracle_vg1 Specify the name for the volume group. The name you specify could be a generic name, as shown, or for a database volume group, it could specify the name of the database that you intend to create.
-B	Create a big VG format Volume Group	Specify this option to create a big VG format volume group. Note: If you are using SMIT, choose yes for this field.
-s <i>PPsize</i>	Physical partition SIZE in megabytes	32 Specify the size of the physical partitions for the database. The sample value shown enables you to include a disk up to 32 GB in size (32 MB * 1016).

Command Option	SMIT Field	Sample Value and Description
-V <i>Majornum</i>	Volume Group MAJOR NUMBER	46 Specify the device major number for the volume group that you identified in Step 7.
-n	Activate volume group AUTOMATICALLY at system restart	Specify this option to prevent the volume group from being activated at system restart. Note: If you are using SMIT, choose no for this field.
-C	Create VG Concurrent Capable	Specify this option to create a concurrent capable volume group. Note: If you are using SMIT, choose yes for this field.
<i>PhysicalVolumes</i>	PHYSICAL VOLUME names	hdisk3 hdisk4 Specify the device names of the disks that you want to add to the volume group.

9. Enter a command similar to the following to vary on the volume group that you created:

```
# /usr/sbin/varyonvg VGname
```

Create Raw Logical Volumes in the New Volume Group

To create the required raw logical volumes in the new volume group, follow these steps:

1. Choose a name for the database that you want to create.

The name that you choose must start with a letter and have no more than four characters, for example, `orcl`.

2. From the following table, identify the logical volumes that you must create:

Number	Partition Size (MB)	Purpose and Sample Logical Volume Name
Raw Logical Volumes for Oracle Database		
1	500	SYSTEM tablespace: <i>dbname_system_raw_500m</i>
1	300 + (Number of instances * 250)	SYSAUX tablespace: <i>dbname_sysaux_raw_800m</i>
Number of instances	500	UNDOTBS n tablespace (One tablespace for each instance, where n is the number of the instances): <i>dbname_undotbsn_raw_500m</i>
1	160	EXAMPLE tablespace: <i>dbname_example_raw_160m</i>
1	120	USERS tablespace: <i>dbname_users_raw_120m</i>

Number	Partition Size (MB)	Purpose and Sample Logical Volume Name
2 * number of instances	120	Two online redo log files for each instance (where <i>n</i> is the number of the instance and <i>m</i> is the log number, 1 or 2): <i>dbname_redon_m_raw_120m</i>
2	110	First and second control files: <i>dbname_control[1 2]_raw_110m</i>
1	250	TEMP tablespace: <i>dbname_temp_raw_250m</i>
1	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwordfile_raw_5m</i>
Raw Logical Volumes for Oracle Cluster Ready Services (CRS)		
1	100	Oracle Cluster Registry: <i>ora_ocr_raw_100m</i> Note: You need to create this raw logical volume only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle cluster registry. If you are upgrading from Oracle9i Release 2, you can continue to use the raw device that you used for the SRVM configuration repository instead of creating this new logical volume.
1	20	Oracle CRS voting disk: <i>ora_vote_raw_20m</i> Note: You need to create this raw logical volume only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle CRS voting disk.

3. To create each required logical volume, enter a command similar to the following (or use `smit mklv`):

```
# /usr/sbin/mklv -y LVname -T 0 -w n -s n -r n VGname NumPPS
```

In this example:

- *LVname* is the name of the logical volume that you want to create
- The `-T 0` option specifies that the device subtype should be `z`, which causes Oracle to use a zero offset when accessing this raw logical volume.
- *VGname* is the name of the volume group where you want to create the logical volume
- *NumPPS* is the number of physical partitions to use

To determine the value to use for *NumPPS*, divide the required size of the logical volume by the size of the physical partition and round the value up to an integer. For example, if the size of the physical partition is 32 MB and you want to create a 500 MB logical volume, you should specify 16 for the *NumPPS* ($500/32 = 15.625$).

The following example shows the command used to create a logical volume for the SYSAUX tablespace of the `test` database in the `oracle_vg1` volume group with a physical partition size of 32 MB ($800/32 = 25$):

```
# /usr/sbin/mkLV -y test_sysaux_800m -T O -w n -s n -r n oracle_vg1 25
```

4. Change the owner, group, and permissions on the character device files associated with the logical volumes that you created, as follows:

Note: The device file associated with the Oracle Cluster Registry must be owned by `root`. All other device files must be owned by the Oracle software owner user (`oracle`).

- Oracle Database files:

```
# chown oracle:dba /dev/rdbname*
# chmod 660 /dev/rdbname*
```

- Oracle CRS files:

```
# chown oracle:dba /dev/rora_vote_raw_20m
# chmod 660 /dev/rora_vote_raw_20m
# chown root:dba /dev/rora_ocr_raw_100m
# chmod 640 /dev/rora_ocr_raw_100m
```

Import the Volume Group on the Other Cluster Nodes

To make the volume group available to all nodes in the cluster, you must import it on each node, as follows:

1. Because the physical volume names might be different on the other nodes, enter the following command to determine the PVID of the physical volumes used by the volume group:

```
# /usr/sbin/lspv
```

2. Note the PVIDs of the physical devices used by the volume group.
3. To vary off the volume group that you want to use, enter commands similar to the following on the node where you created it:

```
# /usr/sbin/varyoffvg VGname
```

4. On each cluster node, complete the following steps:

- a. Enter the following command to determine the physical volume names associated with the PVIDs you noted previously:

```
# /usr/sbin/lspv
```

- b. On each node of the cluster, enter commands similar to the following to import the volume group definitions:

```
# /usr/sbin/importvg -y VGname -V MajorNumber PhysicalVolume
```

In this example, *MajorNumber* is the device major number for the volume group and *PhysicalVolume* is the name of one of the physical volumes in the volume group.

For example, to import the definition of the `oracle_vg1` volume group with device major number 45 on the `hdisk3` and `hdisk4` physical volumes, enter the following command:

```
# /usr/sbin/importvg -y oracle_vg1 -V 45 hdisk3
```

- c. Change the owner, group, and permissions on the character device files associated with the logical volumes you created, as follows:

- Oracle Database files:

```
# chown oracle:dba /dev/rdbname*
# chmod 660 /dev/rdbname*
```

- Oracle CRS files:

```
# chown oracle:dba /dev/rora_vote_raw_20m
# chmod 660 /dev/rora_vote_raw_20m
# chown root:dba /dev/rora_ocr_raw_100m
# chmod 640 /dev/rora_ocr_raw_100m
```

- d. Enter the following command to ensure that the volume group will not be activated by the operating system when the node boots:

```
# /usr/sbin/chvg -a n VGname
```

Activate the Volume Group in Concurrent Mode on All Cluster Nodes

To activate the volume group in concurrent mode on all cluster nodes, enter the following command on each node:

```
# /usr/sbin/varyonvg -c VGname
```

Note: After you install Oracle CRS, the CRS startup scripts use this command to activate the volume group automatically when the node boots.

Create the DBCA Raw Device Mapping File

Note: You must complete this procedure only if you are using raw devices for database files. You do not specify the raw devices for the Oracle CRS files in the DBCA raw device mapping file.

To enable Database Configuration Assistant (DBCA) to identify the appropriate raw device for each database file, you must create a raw device mapping file, as follows:

1. Set the ORACLE_BASE environment variable to specify the Oracle base directory that you identified or created previously:

- Bourne, Bash or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle ; export ORACLE_BASE
```

- C shell:

```
$ setenv ORACLE_BASE /u01/app/oracle
```

2. Create a database file subdirectory under the Oracle base directory and set the appropriate owner, group, and permissions on it:

```
# mkdir -p $ORACLE_BASE/oradata/dbname
# chown -R oracle:oinstall $ORACLE_BASE/oradata
# chmod -R 775 $ORACLE_BASE/oradata
```

In this example, *dbname* is the name of the database that you chose previously.

3. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
4. Enter the following command to create a text file that you can use to create the raw device mapping file:

```
# find /dev -user oracle -name 'r*' -print > dbname_raw.conf
```

5. Edit the *dbname_raw.conf* file in any text editor to create a file similar to the following:

Note: The following example shows a sample mapping file for a two-instance RAC cluster.

```
system=/dev/rddbname_system_raw_500m
sysaux=/dev/rddbname_sysaux_raw_800m
example=/dev/rddbname_example_raw_160m
users=/dev/rddbname_users_raw_120m
temp=/dev/rddbname_temp_raw_250m
undotbs1=/dev/rddbname_undotbs1_raw_500m
undotbs2=/dev/rddbname_undotbs2_raw_500m
redo1_1=/dev/rddbname_redo1_1_raw_120m
redo1_2=/dev/rddbname_redo1_2_raw_120m
redo2_1=/dev/rddbname_redo2_1_raw_120m
redo2_2=/dev/rddbname_redo2_2_raw_120m
control1=/dev/rddbname_control1_raw_110m
control2=/dev/rddbname_control2_raw_110m
spfile=/dev/rddbname_spfile_raw_5m
pwdfile=/dev/rddbname_pwdfile_raw_5m
```

In this example, *dbname* is the name of the database.

Use the following guidelines when creating or editing this file:

- Each line in the file must have the following format:

```
database_object_identifier=logical_volume
```

The logical volume names suggested in this manual include the database object identifier that you must use in this mapping file. For example, in the following logical volume name, *redo1_1* is the database object identifier:

```
/dev/racredo1_1_raw_120m
```

- For a single instance database, the file must specify one automatic undo tablespace data file (*undotbs1*), and at least two redo log files (*redo1_1*, *redo1_2*).
 - For a RAC database, the file must specify one automatic undo tablespace datafile (*undotbsn*) and two redo log files (*redon_1*, *redon_2*) for each instance.
 - Specify at least two control files (*control1*, *control2*).
 - To use manual instead of automatic undo management, specify a single RBS tablespace datafile (*rbs*) instead of the automatic undo management tablespaces.
6. Save the file and note the file name that you specified.

7. When you are configuring the `oracle` user's environment later in this chapter, set the `DBCA_RAW_CONFIG` environment variable to specify the full path to this file.

Verifying the Cluster Software Configuration

Note: The cluster software mentioned in the following sections is not required to install Oracle Real Application Clusters. However, if this software is installed, Oracle Cluster Ready Services (CRS) can integrate with it.

To verify that the HACMP cluster software is configured correctly, follow these steps:

1. If you are using HACMP, verify that it is configured to start at runlevel 2.
2. If you are using GPFS, verify that it is configured to start at runlevel 2. If you are also using HACMP, make sure that GPFS starts after HACMP.
3. To enable the Oracle software owner to use Group Services, make sure that the `oracle` user is a member of the `hagsuser` group:

Note: In this procedure, `oracle` is the user name of the user that installs the Oracle software. If you use different users to install CRS and RAC, both users must belong to the `hagsuser` group.

- a. Create a group named `hagsuser` if it does not exist.
- b. Add the `oracle` user to the `hagsuser` group.
- c. Stop and restart the `groupsvcs` daemon (if it was running before you created the `hagsuser` group).
4. To verify that the HACMP cluster software is configured correctly and is running, enter the following command:

```
# lssrc -ls grpsvcs
```

From the output, verify that the number of providers listed for the group `CLSTRMGR_clusterid` is equal to the number of nodes.

5. To ensure that RAC operates efficiently, you must synchronize the system time on all cluster nodes. Oracle recommends that you use `xntpd` for this purpose. `xntpd` is a complete implementation of the Network Time Protocol (NTP) version 3 standard and is more accurate than `timed`.

To configure `xntpd`, follow these steps on each cluster node:

- a. Enter the following command to create required files, if necessary:

```
# touch /etc/ntp.drift /etc/ntp.trace /etc/ntp.conf
```

- b. Using any text editor, edit the `/etc/ntp.config` file:

```
# vi /etc/ntp.conf
```

- c. Add entries similar to the following to the file:

```
# Sample NTP Configuration file
```

```
# Specify the IP Addresses of three clock server systems.

server ip_address1
server ip_address2
server ip_address3

# Most of the routers are broadcasting NTP time information. If your
# router is broadcasting, then the following line enables xntpd
# to listen for broadcasts.

broadcastclient

# Write clock drift parameters to a file. This enables the system
# clock to quickly synchronize to the true time on restart.

driftfile /etc/ntp.drift
tracefile /etc/ntp.trace
```

- d. To start xntpd, enter the following command

```
# /usr/bin/smitty xntpd
```

then choose **Start Using the xntpd Subsystem**, then choose **BOTH**.

Stop Existing Oracle Processes

Caution: If you are installing additional Oracle Database 10g products in an existing Oracle home, stop **all** processes running in the Oracle home. You must complete this task to enable the Installer to relink certain executables and libraries.

If you choose to create a database during the installation, most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the IPC key value EXTPROC. However, if an existing Oracle Net listener process is using the same port or key value, the Installer can only configure the new listener; it cannot start it. To ensure that the new listener process starts during the installation, you must shut down any existing listeners before starting the Installer.

To determine whether an existing listener process is running and to shut it down if necessary, follow these steps:

1. Switch user to oracle:

```
# su - oracle
```

2. Enter the following command to determine whether a listener process is running and to identify its name and the Oracle home directory in which it is installed:

```
$ ps -ef | grep tnslnsr
```

This command displays information about the Oracle Net listeners running on the system:

```
... oracle_home1/bin/tnslnsr LISTENER -inherit
```

In this example, *oracle_home1* is the Oracle home directory where the listener is installed and LISTENER is the listener name.

3. Set the ORACLE_HOME environment variable to specify the appropriate Oracle home directory for the listener:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=oracle_home1
$ export ORACLE_HOME
```

- C or tcsh shell:

```
$ setenv ORACLE_HOME oracle_home1
```

4. Enter the following command to identify the TCP/IP port number and IPC key value that the listener is using:

```
$ $ORACLE_HOME/bin/lsnrctl status listenername
```

Note: If the listener uses the default name LISTENER, you do not have to specify the listener name in this command.

5. Enter a command similar to the following to stop the listener process:

```
$ORACLE_HOME/bin/lsnrctl stop listenername
```

6. Repeat this procedure to stop all listeners running on this system and on all other cluster nodes.

Configure the oracle User's Environment

You run the Installer from the `oracle` account. However, before you start the Installer you must configure the environment of the `oracle` user. To configure the environment, you must:

- Set the default file mode creation mask (`umask`) to 022 in the shell startup file.
- Set the `DISPLAY` and `ORACLE_BASE` environment variables.

To set the `oracle` user's environment, follow these steps:

1. Start a new terminal session, for example, an X terminal (`xterm`).
2. Enter the following command to ensure that X Window applications can display on this system:

```
$ xhost +
```

3. If you are not already logged in to the system where you want to install the software, log in to that system as the `oracle` user.

4. If you are not logged in as the `oracle` user, switch user to `oracle`:

```
$ su - oracle
```

5. To determine the default shell for the `oracle` user, enter the following command:

```
$ echo $SHELL
```

6. Open the `oracle` user's shell startup file in any text editor:

- Bourne shell (`sh`), Bash shell (`bash`), or Korn shell (`ksh`):

```
$ vi .profile
```

- C shell (csh or tcsh):


```
% vi .login
```
- 7. Enter or edit the following line, specifying a value of 022 for the default file creation mask:


```
umask 022
```
- 8. If the ORACLE_SID, ORACLE_HOME, or ORACLE_BASE environment variables are set in the file, remove the appropriate lines from the file.
- 9. Save the file and exit from the editor.
- 10. To run shell startup script, enter one of the following command:
 - Bourne, Bash, or Korn shell:


```
$ . ./profile
```
 - C shell:


```
% source ./login
```
- 11. If you are not installing the software on the local system, enter a command similar to the following to direct X applications to display on the local system:
 - Bourne, Bash, or Korn shell:


```
$ DISPLAY=local_host:0.0 ; export DISPLAY
```
 - C shell:


```
$ setenv DISPLAY local_host:0.0
```

In this example, *local_host* is the host name or IP address of the system you want to use to display the Installer (your workstation or PC).

- 12. If you determined that the /tmp directory has less than 200 MB of free disk space, identify a file system with at least 200 MB of free space and set the TEMP and TMPDIR environment variables to specify a temporary directory on this file system:
 - a. Use the `df -k` command to identify a suitable file system with sufficient free space.
 - b. If necessary, enter commands similar to the following to create a temporary directory on the file system you identified, and set the appropriate permissions on the directory:


```
$ su - root
# mkdir /mount_point/tmp
# chmod a+wr /mount_point/tmp
# exit
```
 - c. Enter commands similar to the following to set the TEMP and TMPDIR environment variables:
 - * Bourne, Bash, or Korn shell:


```
$ TEMP=/mount_point/tmp
$ TMPDIR=/mount_point/tmp
$ export TEMP TMPDIR
```
 - * C shell:

```
$ setenv TEMP /mount_point/tmp
$ setenv TMPDIR /mount_point/tmp
```

13. Enter commands similar to the following to set the ORACLE_BASE environment variable:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle
$ export ORACLE_BASE
```

- C shell:

```
% setenv ORACLE_BASE /u01/app/oracle
```

In these examples, /u01/app/oracle is the Oracle base directory that you created earlier.

14. If necessary, set the following environment variables. If you must specify several values for an environment variable, for example PATH, separate the values with a colon (:).

Environment Variable	Required By	Sample Setting and Description
DBCA_RAW_CONFIG	Oracle Database installations that use raw devices for database storage.	<code>\$ORACLE_BASE/oradata/dbname/dbname_raw.conf</code> Specifies the location of the raw device mapping file
COBDIR	Pro*COBOL	<code>/opt/lpp/cobol64/cobol</code> Specifies the directory where COBOL is installed on the system. Note: To install Pro*COBOL, you must choose a custom software installation.
PATH	Pro*COBOL	<code>\$COBDIR/bin:\$PATH</code> Specifies the directory where the COBOL compiler executable is located.

15. Enter the following command to ensure that the ORACLE_HOME and TNS_ADMIN environment variable are not set:

```
$ unset ORACLE_HOME
$ unset TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value it specifies as the default path for the Oracle home directory. However, if you set the ORACLE_BASE environment variable, Oracle recommends that you unset this environment variable and choose the default path suggested by the Installer.

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

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

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

Pre-Installation Tasks for Installing RAC on HP-UX Systems

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer. It includes information about the following tasks:

- Log In to the System as root
- Check the Hardware Requirements
- Check the Network Requirements
- Check the Software Requirements
- Create Required UNIX Groups and User
- Configure Kernel Parameters
- Identify Required Software Directories
- Identify or Create an Oracle Base Directory
- Configure Disk Storage for Oracle CRS and Database Files
- Configure Disks for Automatic Storage Management
- Configure Raw Logical Volumes
- Verify the Cluster Software Configuration
- Stop Existing Oracle Processes
- Configure the oracle User's Environment

Log In to the System as root

Before you install the Oracle software, you must complete several tasks as the `root` user. To log in as the `root` user, complete one of the following procedures:

Note: Unless you intend to complete a silent installation, you must install the software from an X Window System workstation, an X terminal, or a PC or other system with X server software installed.

For more information about non-interactive and silent installations, see the appendices in *Oracle Database Installation Guide for UNIX Systems*.

- If you are installing the software from an X Window System workstation or X terminal:
 1. Start a local terminal session, for example, an X terminal (`xterm`).
 2. If you are not installing the software on the local system, enter the following command to enable remote hosts to display X applications on the local X server:

```
$ xhost +
```
 3. If you want to install the software on a remote system, enter a command similar to the following to connect to that system:

```
$ telnet remote_host
```
 4. If you are not logged in as the `root` user, enter the following command to switch user to `root`:

```
$ su - root
password:
#
```
- If you are installing the software from a PC or other system with X server software installed:

Note: If necessary, see your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

1. Start the X server software.
2. Configure the security settings of the X server software to permit remote hosts to display X applications on the local system.
3. Connect to the remote system where you want to install the software and start a terminal session on that system, for example, an X terminal (`xterm`).
4. If you are not logged in as the `root` user on the remote system, enter the following command to switch user to `root`:

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

Check the Hardware Requirements

The system must meet the following minimum hardware requirements:

- 512 MB of physical RAM
- 1 GB of swap space (or twice the size of RAM)
On systems with 2 GB or more of RAM, the swap space can be between one and two times the size of RAM.
- 400 MB of disk space in the `/tmp` directory
- Up to 4 GB of disk space for the Oracle software, depending on the installation type
- Between 0.5 GB and 3 GB of disk space for the Oracle software, depending on the installation type
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirements for databases that use Automatic Storage Management (ASM) or raw device storage is described later in this chapter.

Additional disk space, either on a file system or in an ASM disk group, is required for the flash recovery area if you choose to configure automated backups.

To ensure that the system meets these requirements, follow these steps:

1. To determine the physical RAM size, enter the following command:

```
# /usr/sbin/dmmsg | grep "Physical:"
```

If the size of the physical RAM installed in the system is less than the required size, you must install more memory before continuing.

2. To determine the size of the configured swap space, enter the following command:

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

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

3. To determine the amount of disk space available in the `/tmp` directory, enter the following command:

```
# bdf /tmp
```

If there is less than 400 MB of disk space available in the `/tmp` directory, complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to achieve the required disk space.
- Set the `TEMP` and `TMPDIR` environment variables when setting the `oracle` user's environment (described later).
- Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.

4. To determine the amount of free disk space on the system, enter the following command:

```
# bdf
```

5. To determine whether the system architecture can run the software, enter the command:

```
# /bin/getconf KERNEL_BITS
```

The expected output is the value 64.

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

Check the Network Requirements

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Real Application Clusters installation.

Note: For the most up-to-date information about supported network protocols and hardware for RAC installations, see the Certify pages on the Oracle*MetaLink* Web site:

<http://metalink.oracle.com>

Network Hardware Requirements

Each node in the cluster must meet the following requirements:

- Each node must have at least two network adapters; one for the public network interface and one for the private network interface (the interconnect).
- The interface names associated with the network adapters for each network must be the same on all nodes.
- For increased reliability, you can configure redundant public and private network adapters for each node.
- For the public network, each network adapter must support TCP/IP.
- For the private network, Oracle supports the following interconnect protocols and hardware:
 - User datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended)

Note: UDP is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS.

- Hyper Messaging protocol (HMP) using HyperFabric adapters and switches

IP Address Requirements

Before starting the installation, you must identify or obtain the following IP addresses for each node:

- An IP address and an associated host name registered in the domain name service (DNS) for each public network interface
- One unused virtual IP address and an associated virtual host name registered in DNS that you will configure for the primary public network interface

The virtual IP address must be in the same subnet as the associated public interface. After installation, you can configure clients to use the virtual host name or IP address. If a node fails, its virtual IP address fails over to another node.

- A private IP address and optional host name for each private interface
- Oracle recommends that you use private network IP addresses for these interfaces, for example: 10.*.* or 192.168.*.*. You can use the `/etc/hosts` file on each node to associate private host names with private IP addresses.

For example, if each node has two public and two private interfaces, you might have the following host names and IP addresses on one of the nodes (`rac1`) and similar host names and IP addresses on the other nodes:

Host Name	Type	IP Address	Registered In
<code>rac1.mydomain.com</code>	Public	143.47.43.100	DNS
<code>rac1-2.mydomain.com</code>	Public	143.46.51.101	DNS
<code>rac1-vip.mydomain.com</code>	Virtual	143.46.43.104	DNS
<code>rac1-priv1</code>	Private	10.0.0.1	<code>/etc/hosts</code>
<code>rac1-priv2</code>	Private	10.0.0.2	<code>/etc/hosts</code>

Checking the Network Requirements

To verify that each node meets the requirements, follow these steps:

1. If necessary, install the network adapters for the public and private networks and configure them with either public or private IP addresses.
2. Register the host names and IP addresses for the public network interfaces in DNS.
3. For each node, register one virtual host name and IP address in DNS.
4. For each private interface on every node, add a line similar to the following to the `/etc/hosts` file on all nodes, specifying the private IP address and associated private host name:

```
10.0.0.1    rac1-priv1
```

5. To identify the interface name and associated IP address for every network adapter, enter the following command:

```
# /usr/bin/netstat -in
```

From the output, identify the interface name and IP address for all network adapters that you want to specify as public or private network interfaces.

Note: When you install Oracle CRS and RAC, you will require this information.

Check the Software Requirements

See the following sections for information about checking the software requirements:

- [Check for Required Software](#) on page 3-6
- [Check for Required Patches](#) on page 3-7
- [Create Required Symbolic Links](#) on page 3-9

Note: The Oracle Universal Installer performs checks on your system to verify that it meets the requirements listed for HP-UX. To ensure that these checks pass, verify the requirements before you start the Installer.

Check for Required Software

Depending on the products that you intend to install, verify that the following software is installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	Operating system version: HP-UX 11i (11.11) PA-RISC
Oracle Spatial	HP-UX Developer's Toolkit - X11, Motif, and Imake: X11MotifDevKit Note: This software is required only to build the sample programs.
Real Application Clusters	<ul style="list-style-type: none"> ■ HP Serviceguard A11.15.00 or higher ■ HP Serviceguard Extension for RAC A11.15.00 or higher ■ HyperFabric software - Required if you want to use Hyper Messaging protocol as the cluster interconnect ■ .01 or higher Note: HP Serviceguard is required only if you want to use shared logical volumes for CRS or database files.
Oracle Messaging Gateways	IBM WebSphere MQ V5.3, client and server: MQSERIES.MQM-CL-HPUX MQSERIES.MQM-SERVER TIBCO Rendezvous 7.2
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	HP C/ANSI C Compiler (B.11.11.08 or higher): C-ANSI-C HP aC++ Compiler (A.03.50 or higher): ACXX
Oracle JDBC/OCI Drivers	You can use the following optional Java SDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation: <ul style="list-style-type: none"> ■ Java SDK 1.3.1.02 with the JNDI extension ■ Java SDK 1.2.2.09 Note: Java SDK 1.4.2.00 is installed with this release.

Installation Type or Product	Requirement
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE	DCE (B.11.11.10 or higher): DCE-C-Tools DCE-CDS-Server DCE-Core DCE-CoreAdmin DCE-CoreTools DCE-SEC-Server
Oracle Transparent Gateway for Sybase	Either of the following: <ul style="list-style-type: none"> ■ Sybase Adaptive Server Enterprise V12 or 12.5 ■ Sybase Open Client V12 or 12.5 Open Client is required if the Sybase server is not on the local node.
Oracle Transparent Gateway for Informix	One of the following: <ul style="list-style-type: none"> ■ Informix Dynamic Server v7.3, 7.4, or 9.* ■ Informix ESQ/C v7.3, 7.4, or 9.* ESQ/C is required if the Informix server is not on the local node. Note: This version of the gateway does not support Informix SE.
Oracle Transparent Gateway for Ingress	Ingres II version 2.0, 2.5, or 2.6
Oracle Transparent Gateway for Teradata	NCR Teradata ODBC Driver v02.04.00.00

To ensure that the system meets these requirements, follow these steps:

1. To determine which version of HP-UX is installed, enter the following command:

```
# uname -a
HP-UX hostname B.11.11 U 9000/800 109444686 unlimited-user license
```

In this example, the version of HP-UX 11i is 11.11.

2. To determine whether a bundle, product, or file set is installed, enter a command similar to the following, where *level* is *bundle*, *product*, or *fileset*:

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

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

Check for Required Patches

Depending on the products that you intend to install, verify that the following patches are installed on the system. The procedure following the table describes how to check these requirements.

Note: There may be more recent versions of the patches listed installed on the system. If a listed patch is not installed, determine whether a more recent version is installed before installing the version listed.

Installation Type or Product	Requirement
All installations	Quality Pack bundle: HP-UX 11i Quality Pack (GOLDQPK11i), June 2003 or later: GOLDAPPS11i GOLDBASE11i
All installations	Patches for HP-UX 11i (11.11): <ul style="list-style-type: none"> ■ PHCO_28123: cumulative SAM patch ■ PHKL_29198: Psets Enablement Patch; top(1) ■ PHNE_28476: Cumulative STREAMS Patch ■ PHNE_28923: LAN product cumulative patch ■ PHSS_28871: ld(1) and linker tools cumulative patch ■ PHSS_28880: HP aC++ -AA runtime libraries (aCC A.03.50)
All installations	Patches for JDK on HP-UX 11i (11.11): <ul style="list-style-type: none"> ■ PHCO_26331: mountall cumulative patch, Dev IDs enabler ■ PHCO_29109: Pthread enhancement and fixes ■ PHKL_25468: eventport (/dev/poll) pseudo driver ■ PHKL_25842: Thread Abort ■ PHKL_25993: thread nostop for NFS, rlimit, Ufalloc fix ■ PHKL_25994: Thread NOSTOP, Psets Enablement, Ufalloc ■ PHKL_25995: ufalloc;VxFS3.5;SPP fragmentation ■ PHKL_26468: Shared synchronization performance support ■ PHKL_28489: copyin EFAULT, LDCD access type <p>Note: See the following Web site for information about additional patches that might be required by JDK 1.4.2: http://www.hp.com/products1/unix/java/patches/index.html</p>
Real Application Clusters	Patches for HP-UX 11i (11.11): PHSS_29096
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Patches for HP-UX 11i (11.11): <ul style="list-style-type: none"> ■ PHSS_29484: ANSI C compiler B.11.11.08 cumulative patch ■ PHSS_29485: +O4/PBO Compiler B.11.11.08 cumulative patch
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE	Patches for HP-UX 11i (11.11): <ul style="list-style-type: none"> ■ PHSS_28386: HP DCE/9000 1.8 DCE Client IPv6 patch ■ PHSS_28387: HP DCE/9000 1.8 Server/DevTools cum. patch ■ PHSS_28388: HP DCE/9000 1.8 Integrated Login cum. patch ■ PHSS_29669: HP DCE/9000 1.8 DCE Client IPv6 patch
Oracle Messaging Gateway	Corrective service diskettes (CSDs) for WebSphere MQ: <ul style="list-style-type: none"> ■ CSD05 or later for WebSphere MQ V5.3

To ensure that the system meets these requirements, follow these steps:

1. To determine whether the HP-UX 11i Quality Pack is installed, enter the following command:

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

If the Quality Pack is not installed, or if the date is before June 2003, download the latest Quality Pack from the following Web site and install it:

```
http://www.software.hp.com/SUPPORT_PLUS/gpk.html
```

2. To determine whether a patch is installed, enter a command similar to the following:

```
# /usr/sbin/swlist -l patch | grep PHSS_28880
```

Alternatively, to list all installed patches, enter the following command:

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

If a required patch is not installed, download it from the following Web site and install it:

```
http://itresourcecenter.hp.com
```

If the Web site shows a more recent version of the patch, download and install that version.

3. If you require a CSD for WebSphere MQ, see the following Web site for download and installation information:

```
http://www.ibm.com/software/integration/mqfamily/support/summary/hpx.html
```

Create Required Symbolic Links

Note: This task is required only if the Motif 2.1 Development Environment package (X11MotifDevKit.MOTIF21-PRG) is not installed.

To enable you to successfully relink Oracle products after installing this software, enter the following commands to create required X library symbolic links in the `/usr/lib` directory:

```
# cd /usr/lib
# ln -s libX11.3 libX11.sl
# ln -s libXIE.2 libXIE.sl
# ln -s libXext.3 libXext.sl
# ln -s libXhp11.3 libXhp11.sl
# ln -s libXi.3 libXi.sl
# ln -s libXm.4 libXm.sl
# ln -s libXp.2 libXp.sl
# ln -s libXt.3 libXt.sl
# ln -s libXtst.2 libXtst.sl
```

Create Required UNIX Groups and User

Depending on whether this is the first time Oracle software is being installed on this system and on the products that you are installing, you may need to create several UNIX groups and a UNIX user.

The following UNIX groups and user are required if you are installing Oracle Database:

- The OSDBA group (`dba`)

You must create this group the first time you install Oracle Database software on the system. It identifies UNIX users that have database administrative privileges (the SYSDBA privilege). The default name for this group is `dba`.

If you want to specify a group name other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of this group. In this case, the Installer prompts you to specify the name of this group.

- The OSOPER group (`oper`)

This is an optional group. Create this group if you want a separate group of UNIX users to have a limited set of database administrative privileges (the SYSOPER privilege). By default, members of the OSDBA group also have the SYSOPER privilege.

If you want to specify a separate OSOPER group, other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of the `dba` group. In this case, the Installer prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

- An unprivileged user

You must create a new unprivileged user. The user that you create must own the external jobs (`ext job`) executable after the installation.

The following UNIX group and user are required for all installations:

- The Oracle Inventory group (`oinstall`)

You must create this group the first time you install Oracle software on the system. The usual name chosen for this group is `oinstall`. This group owns the Oracle inventory, which is a catalog of all Oracle software installed on the system.

Note: If Oracle software is already installed on the system, the existing Oracle Inventory group must be the primary group of the UNIX user that you use to install new Oracle software. The following sections describe how to identify an existing Oracle Inventory group.

- The Oracle software owner user (`oracle`)

You must create this user the first time you install Oracle software on the system. This user owns all of the software installed during the installation. The usual name chosen for this user is `oracle`. This user must have the Oracle Inventory group as its primary group. It must also have the OSDBA and OSOPER groups as a secondary groups.

A single Oracle Inventory group is required for all installations of Oracle software on the system. After the first installation of Oracle software, you must use the same Oracle Inventory group for all subsequent Oracle software installations on that system. However, you can choose to create different Oracle software owner users, OSDBA groups, and OSOPER groups (other than `oracle`, `dba`, and `oper`) for separate installations. By using different groups for different installations, members of these different groups have DBA privileges only on the associated databases rather than on all databases on the system.

See Also: The *Oracle Database Administrator's Reference for UNIX Systems* and the *Oracle Database Administrator's Guide* contain more information about the OSDBA and OSOPER groups and the SYSDBA and SYSOPER privileges.

The following sections describe how to create the required UNIX user and groups.

Note: The following sections describe how to create local users and groups. As an alternative to creating local users and groups, you could create the appropriate users and groups in a directory service, for example, Network Information Services (NIS). For information about using directory services, contact your system administrator or see your operating system documentation.

Creating the Oracle Inventory Group

You must create the Oracle Inventory group if it does not already exist. The following subsections describe how to determine the Oracle Inventory group name, if it exists, and how to create it if necessary.

Determine Whether the Oracle Inventory Group Exists

When you install Oracle software on the system for the first time, the Installer creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group and the path of the Oracle Inventory directory. To determine whether the Oracle Inventory group exists, enter the following command:

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

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inst_group` parameter shows the name of the Oracle Inventory group (`oinstall`).

Create the Oracle Inventory Group

If the `oraInst.loc` file does not exist, create the Oracle Inventory group, `oinstall`, by entering the following command:

```
# /usr/sbin/groupadd oinstall
```

Creating the OSDBA Group

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system
- An OSBDA group exists, but you want to give a different group of UNIX users database administrative privileges in a new Oracle installation

If the OSDBA group does not exist or if you require a new OSDBA group, create it by entering the following command, using the group name `dba` unless a group with that name already exists:

```
# /usr/sbin/groupadd dba
```

Creating an OSOPER Group (Optional)

Create an OSOPER group only if you want to identify a group of UNIX users with a limited set of database administrative privileges (SYSOPER operator privileges). For most installations, it is sufficient to create only the OSDBA group. If you want to use an OSOPER group, you must create it in the following circumstances:

- If an OSOPER group does not exist, for example, if this is the first installation of Oracle Database software on the system
- If an OSOPER group exists, but you want to give a different group of UNIX users database operator privileges in a new Oracle installation

If you require a new OSOPER group, create it by entering the following command, using the group name `oper` unless a group with that name already exists:

```
# /usr/sbin/groupadd oper
```

Creating the Oracle Software Owner User

You must create an Oracle software owner user in the following circumstances:

- If an Oracle software owner user does not exist, for example, if this is the first installation of Oracle software on the system
- If an Oracle software owner user exists, but you want to use a different UNIX user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation

Determine Whether an Existing Oracle Software Owner User Exists

To determine whether an Oracle software owner user named `oracle` exists, enter the following command:

```
# id oracle
```

If the `oracle` user exists, the output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

If the user exists, determine whether you want to use the existing user or create a new user. If you want to use the existing user, ensure that the user's primary group is the Oracle Inventory group and that it is a member of the appropriate OSDBA and OSOPER groups. See one of the following sections for more information:

Note: If necessary, contact your system administrator before using or modifying an existing user.

- If you want to use the existing Oracle software owner user, and the user's primary group is the Oracle Inventory group, see ["Creating an Unprivileged User"](#) on page 3-13
- To modify an existing user, see the ["Modify an Existing Oracle Software Owner User"](#) section on page 3-13.
- To create a new user, see the following section.

Create a New Oracle Software Owner User

If the Oracle software owner user does not exist or if you require a new Oracle software owner user, create it as follows, using the user name `oracle` unless a user with that name already exists.

1. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example, `oinstall`
 - The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba , oper`
2. Set the password of the `oracle` user:

```
# passwd oracle
```

See ["Creating an Unprivileged User"](#) on page 3-13 to continue:

Modify an Existing Oracle Software Owner User

If the `oracle` user exists, but its primary group is not `oinstall` or it is not a member of the appropriate OSDBA or OSOPER groups, you can modify it by entering a command similar to the following, specifying the primary group using the `-g` option and any required secondary groups using the `-G` option: as follows:

```
# /usr/sbin/usermod -g oinstall -G dba[,oper] oracle
```

Creating an Unprivileged User

You must create a new unprivileged local user to own the `ext job` executable. You can use any user name for this user, but this document uses the name `ext job`. To create the `ext job` user, enter the following command on all cluster nodes:

```
# /usr/sbin/useradd extjob
```

Create Identical Users and Groups on Other Cluster Nodes

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

The Oracle software owner user and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes. To create these identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, then create the user and groups with the same name and ID on the other cluster nodes.

Identifying the User and Group IDs

To determine the user ID (UID) of the Oracle software owner user and the group IDs (GID) of the Oracle Inventory, OSDBA, and OSOPER groups, follow these steps:

1. Enter the following command:

```
# id oracle
```

The output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

2. From the output, identify the UID for the `oracle` user and the GIDs for the groups to which it belongs.

Create the User and Groups on the Other Cluster Nodes

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

1. Log in to the next cluster node as `root`.
2. Enter commands similar to the following to create the `oinstall` and `dba` groups, and if required, the `oper` group. Use the `-g` option to specify the correct GID for each group.

```
# /usr/sbin/groupadd -g 200 oinstall
# /usr/sbin/groupadd -g 201 dba
# /usr/sbin/groupadd -g 202 oper
```

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

3. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -u 200 -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-u` option specifies the user ID, which must be the user ID that you identified in the previous subsection.
- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`
- The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba ,oper`

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

4. Set the password of the `oracle` user:

```
# passwd oracle
```

Set Up User Equivalence for rsh and rcp on All Cluster Nodes

Note: This section describes how to set up user equivalence for `rcp`, which the Installer uses when copying Oracle software to the other cluster nodes. If you prefer, you can configure the Secure Shell (SSH) tool suite, so that the Installer uses `scp` instead of `rcp`. See the SSH documentation for information about setting up user equivalence for `scp`.

Before you install and use Oracle Real Application clusters, you must set up user equivalence on all cluster nodes, as follows:

1. On all cluster nodes, edit the `/etc/hosts.equiv` file in any text editor and add entries similar to the following for all nodes in the cluster:

```
node1-public oracle
node1-public.domain oracle
node2-public oracle
node2-public.domain oracle
.
.
.
```

In this example:

- `noden_public` is the host name associated with the primary public network interface
- `domain` is the domain name of the system, for example `mydomain.com`

Note: Alternatively, you can create a `.rhosts` file using the same format in the `oracle` user's home directory on all nodes in the cluster.

2. To test user equivalence on all cluster nodes, including the first node, enter commands similar to the following from the node where you will run the Installer:

```
# su oracle
$ remsh nodename1 date
$ remsh nodename2 date
.
.
.
$ exit
```

Each command should return only the date setting on the specified node. If a node prompts for a password or displays any other text, then user equivalence is not set up correctly on that node. In this case, verify that the `/etc/hosts.equiv` (or `~oracle/.rhosts`) file has the correct entries.

Grant Privileges to the OSDBA Group

To enhance the performance of the software, you must grant several system privileges to the OSDBA group by completing the following steps:

1. Using any text editor, open the `/etc/privgroup` file, or create it if necessary.

2. Add or edit the following line, which begins with the name of the OSDBA group, specifying the privileges MLOCK, RTPRIO, and RTSCHED, that you want to grant to this group every time the system reboots:

```
dba MLOCK RTPRIO RTSCHED
```

3. Save the file and exit from the text editor.
4. Enter the following command to grant the privileges to the OSDBA group:

```
# /usr/sbin/setprivgrp -f /etc/privgroup
```

5. Enter the following command to verify that the privileges are set correctly:

```
# /usr/sbin/getprivgrp dba
```

6. Repeat this procedure on all of the other cluster nodes.

Configure Kernel Parameters

Notes: The kernel parameter and shell limit values shown in this section are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

You must set the kernel parameters on all cluster nodes.

Verify that the kernel parameters shown in the following table are set either to the formula shown, or to values greater than or equal to the recommended value shown. The procedure following the table describes how to verify and set the values.

Parameter	Recommended Formula or Value
ksi_alloc_max	(nproc*8)
max_thread_proc	256
maxdsiz	1073741824 (1 GB)
maxdsiz_64bit	2147483648 (2 GB)
maxssiz	134217728 (128 MB)
maxssiz_64bit	1073741824 (1 GB)
maxswapchunks	16384
maxuprc	((nproc*9)/10)
msgmap	(2+msgmni)
msgmni	4096
msgseg	32767
msgtql	4096
ncsize	(ninode+vx_ncsize)
nfile	(15*nproc+2048)
nflocks	4096

Parameter	Recommended Formula or Value
ninode	$(8 * nproc + 2048)$
nkthread	$((nproc * 7) / 4) + 16$
nproc	4096
semmap	$(semnmi + 2)$
semmni	4096
semmns	$(semnmi * 2)$
semmnu	$(nproc - 4)$
semvmx	32767
shmmax	The size of physical memory (0X40000000) or 1073741824, whichever is greater. Note: To avoid performance degradation, the value should be greater than or equal to the size of the SGA. If you are using Hyper Messaging protocol as the cluster interconnect, the value of this parameter must be less than 4 GB.
shmmni	512
shmseg	120
vps_ceiling	64

Note: If the current value for any parameter is higher than the value listed in this table, do not change the value of that parameter.

To view the current value or formula specified for these kernel parameters, and to change them if necessary, follow these steps:

- Optionally, set the DISPLAY environment variable to specify the display of the local system:
 - Bourne, Bash, or Korn shell:

```
$ DISPLAY=local_host:0.0 ; export DISPLAY
```
 - C shell:

```
$ setenv DISPLAY local_host:0.0
```
- Start System Administration Manager (SAM):

```
# /usr/sbin/sam
```
- Choose the **Kernel Configuration** area, then choose the **Configurable Parameters** area.
- Check the value or formula specified for each of these parameters and, if necessary, modify that value or formula.
If necessary, see the SAM online help for more information about completing this step.
- Exit from SAM.
If you modified the value specified for any static parameter, you must rebuild the kernel and reboot the system.
- If necessary, when the system restarts, log in and switch user to root.

7. Complete this procedure on all other cluster nodes.

Identify Required Software Directories

You must identify or create four directories for the Oracle software, as follows:

- Oracle base directory
- Oracle Inventory directory
- CRS home directory
- Oracle home directory

The following subsections describe the requirements for these directories.

Oracle Base Directory

The Oracle base directory acts as a top-level directory for Oracle software installations. It is analogous to the `C:\Oracle` directory used for Oracle software installations on Windows systems. On UNIX systems, the Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to the following for the Oracle base directory:

```
/mount_point/app/oracle_sw_owner
```

In this example:

- *mount_point* is the mount point directory for the file system that will contain the Oracle software.

The examples in this guide use `/u01` for the mount point directory. However, you could choose another mount point directory, `/oracle` or `/opt/oracle` for example.

- *oracle_sw_owner* is the UNIX user name of the Oracle software owner, for example `oracle`.

You can use the same Oracle base directory for more than one installation or you can create separate Oracle base directories for different installations. If different UNIX users install Oracle software on the same system, each user must create a separate Oracle base directory. The following example Oracle base directories could all exist on the same system:

```
/u01/app/oracle  
/u01/app/orauser  
/opt/oracle/app/oracle
```

The following sections describe how to identify existing Oracle base directories that might be suitable for your installation and how to create a new Oracle base directory if necessary.

Regardless of whether you create a new Oracle base directory or decide to use an existing one, you must set the `ORACLE_BASE` environment variable to specify the full path to this directory.

Note: The Oracle base directory must be on a local file system.

Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system. It is required by, and shared by, all Oracle software installations on a single system. The first time you install Oracle software on a system, the Installer prompts you to specify the path to this directory. If you are installing the software on a local file system, Oracle recommends that you choose the following path:

```
oracle_base/oraInventory
```

If the Oracle base directory is on a cluster file system, you must specify a path for the Oracle Inventory directory on a local file system, to enable all of the nodes to have separate inventories.

The Installer creates the directory that you specify and sets the correct owner, group, and permissions on it. You do not need to create it.

Note: All Oracle software installations rely on this directory. Make sure that you back it up regularly.

Do not delete this directory unless you have completely removed all Oracle software from the system.

CRS Home Directory

The CRS home directory is the directory where you choose to install the software for Oracle Cluster Ready Services. You must install CRS in a separate home directory. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
oracle_base/product/10.1.0/crs
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Oracle Home Directory

The Oracle home directory is the directory where you choose to install the software for a particular Oracle product. You must install different Oracle products, or different releases of the same Oracle product, in separate Oracle home directories. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the Oracle home directory:

```
oracle_base/product/10.1.0/db_1
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Identify or Create an Oracle Base Directory

Before starting the installation, you must either identify an existing Oracle base directory or if required, create a new one. This section contains information about the following:

- [Identifying an Existing Oracle Base Directory](#)
- [Creating a New Oracle Base Directory](#)

Note: You can choose to create a new Oracle base directory, even if other Oracle base directories exist on the system.

Identifying an Existing Oracle Base Directory

Existing Oracle base directories might not have paths that comply with OFA guidelines. However, if you identify an existing Oracle Inventory directory or existing Oracle home directories, you can usually identify the Oracle base directories, as follows:

- Identify an existing Oracle Inventory directory

Enter the following command on all cluster nodes to view the contents of the `oraInst.loc` file:

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

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inventory_loc` parameter identifies the Oracle Inventory directory (`oraInventory`) on that system. The parent directory of the `oraInventory` directory is typically an Oracle base directory. In the previous example, `/u01/app/oracle` is an Oracle base directory.

- Identify existing Oracle home directories

Enter the following command on all cluster nodes to view the contents of the `oratab` file:

```
# more /etc/oratab
```

If the `oratab` file exists, it contains lines similar to the following:

```
*:/u03/app/oracle/product/10.1.0/db_1:N
*/opt/orauser/infra_904:N
*/oracle/9.2.0:N
```

The directory paths specified on each line identify Oracle home directories. Directory paths that end with the user name of the Oracle software owner that you want to use are valid choices for an Oracle base directory. If you intend to use the `oracle` user to install the software, you could choose one of the following directories from the previous example:

```
/u03/app/oracle
/oracle
```

Note: If possible, choose a directory path similar to the first (/u03/app/oracle). This path complies with the OFA guidelines.

Before deciding to use an existing Oracle base directory for this installation, make sure that it satisfies the following conditions:

- It should not be on the same file system as the operating system.
- It must have up to 4 GB of free disk space. Use the following command to determine the free disk space on the file system where the Oracle base directory is located:

```
# bdf oracle_base_path
```

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the "[Creating a New Oracle Base Directory](#)" section on page 3-21.
- If an Oracle base directory exists and you want to use it, see the "[Configure Disk Storage for Oracle CRS and Database Files](#)" section on page 3-22.

When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify the directory you chose.

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the following section.

Creating a New Oracle Base Directory

Before you create a new Oracle base directory, you must identify an appropriate file system with at least 4 GB of free disk space.

To identify an appropriate file system, follow these steps:

1. Use the `bdf` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has appropriate free space.

Note: The file system must be a local file system.

The path to the Oracle base directory must be the same on all nodes.

3. Note the name of the mount point directory for the file system that you identified.

To create the Oracle base directory and specify the correct owner, group, and permissions for it, follow these steps:

1. Enter commands similar to the following to create the recommended subdirectories in the mount point directory that you identified and set the appropriate owner, group, and permissions on them:

```
# mkdir -p /mount_point/app/oracle_sw_owner
# chown -R oracle:oinstall /mount_point/app
# chmod -R 775 /mount_point/app
```

If the mount point you identified is /u01 and `oracle` is the user name of the Oracle software owner, the recommended Oracle base directory path is as follows:

```
/u01/app/oracle
```

2. If necessary, repeat the commands listed in the previous step to create the same directory on the other cluster nodes.
3. When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify this directory.

Configure Disk Storage for Oracle CRS and Database Files

This section describes the storage options available to you for storing Oracle Cluster Ready Services files, Oracle Database Files, and optionally, Oracle Database recovery files. After you choose the storage method that you want to use for each file type, see the following sections for information about configuring the required storage.

Note: You do not have to use the same storage option for each type of file.

Choosing a Storage Option for Oracle CRS Files

Before you install Oracle Cluster Ready Services, you must choose the storage option that you want to use for the Oracle Cluster Registry (100 MB) and CRS voting disk (20 MB). You cannot use Automatic Storage Management to store these files, because they must be accessible before any Oracle instance starts.

Apart from ASM, you can use any of the storage options listed in the next section for the Oracle CRS files.

Note: If you are not using HP Serviceguard, you must use raw partitions to store the Oracle CRS files. You cannot use shared raw logical volumes to store these files without HP Serviceguard.

Choosing a Storage Option for Oracle Database Files

If you want to create a database during the installation, you must choose one of the following storage options for the database files:

- Automatic Storage Management
- Shared logical volumes using HP Serviceguard Extensions for RAC and LVM
- Raw partitions

Note: For the most up-to-date information about supported storage options for RAC installations, see the Certify pages on the Oracle*MetaLink* Web site:

<http://metalink.oracle.com>

Planning a Storage Option for Oracle Database Recovery Files

Additionally, if you want to enable automated backups during the installation, you must use Automatic Storage Management for recovery files (the flash recovery area).

Configuring Disk Storage

For information about how to configure disk storage before you start the installation, see one of the following sections depending on your choice:

- To use ASM for database or recovery file storage, see the "[Configure Disks for Automatic Storage Management](#)" section on page 3-23.
- To use raw devices (partitions or logical volumes) for Oracle CRS or database file storage, see the "[Configure Raw Logical Volumes](#)" section on page 3-27.

Configure Disks for Automatic Storage Management

This section describes how to configure disks for use with ASM. Before you configure the disks, you must determine the number of disks and the amount of free disk space that you require. The following sections describe how to identify the requirements and configure the disks:

- [Identifying Storage Requirements for ASM](#)
- [Using an Existing ASM Disk Group](#)
- [Configuring Disks for ASM](#)

Note: Although this section refers to disks, you can also use zero-padded files on a certified NAS storage device in an ASM disk group. See the appendices in *Oracle Database Installation Guide for UNIX Systems* for information about creating and configuring NAS-based files for use in an ASM disk group.

Identifying Storage Requirements for ASM

To identify the storage requirements for using ASM, you must determine how many devices and the amount of free disk space that you require. To complete this task, follow these steps:

1. Determine whether you want to use ASM for Oracle database files (datafiles), recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. One can use the file system, while the other uses ASM.

If you enable automated backups during the installation, you can choose ASM as the storage mechanism for recovery files by specifying an ASM disk group for the flash recovery area. Depending how you choose to create a database during the installation, you have the following options:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to use the same ASM disk group for datafiles and recovery files, or you can choose to use different disk groups for each file type.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must use the same ASM disk group for database files and recovery files.

2. Choose the ASM redundancy level that you want to use for the ASM disk group.

The redundancy level that you choose for the ASM disk group determines how ASM mirrors files in the disk group and determines the number of disks and amount of disk space that you require, as follows:

- 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 ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use only RAID or similar devices that provide their own data protection mechanisms as disk devices in this type of disk group.

- Normal redundancy

In a normal redundancy disk group, ASM uses two-way mirroring by default, to increase performance and reliability. 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 most installations, Oracle recommends that you use normal redundancy disk groups.

- High redundancy

In a high redundancy disk group, 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.

While high redundancy disk groups do provide a high level of data protection, you must consider the higher cost of additional storage devices before deciding to use this redundancy level.

3. Determine the total amount of disk space that you require for the datafiles and recovery files.

Use the following table to determine the minimum number of disks and the minimum disk space requirements for the installation:

Redundancy Level	Minimum Number of Disks	Datafiles	Recovery Files	Both File Types
External	1	1.15 GB	2.3 GB	3.45 GB
Normal	2	2.3 GB	4.6 GB	6.9 GB
High	3	3.45 GB	6.9 GB	10.35 GB

You must also add additional disk space for the ASM metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

$$15 + (2 * \text{number_of_disks}) + (126 * \text{number_of_ASM_instances})$$

For example, for a four-node RAC installation, using three disks in a high redundancy disk group, you require an additional 525 MB of disk space:

$$(15 + (2 * 3) + (126 * 4)) = 525$$

If an existing ASM instance exists on the system, 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 installation.

The following section describes how to identify existing disk groups and determine the free disk space that they contain.

4. Optionally identify failure groups for the ASM disk group devices.

Note: You need to complete this step only if you intend to use an installation method that runs DBCA in interactive mode, for example, if you intend to choose the Custom installation type or the Advanced database configuration option. Other installation types do not enable you to specify failure groups.

If you intend to use a normal or high redundancy disk group, 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 SCSI controller, the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To avoid failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

5. If you are sure that a suitable disk group does not exist on the system, 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 ASM disk group should be the same size and have the same performance characteristics.
 - Do not specify more than one partition on a single physical disk as a disk group device. 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 ASM disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing ASM from optimizing I/O across the physical devices.

For information about completing this task, see the "[Configuring Disks for ASM](#)" section on page 3-27.

Using an Existing ASM Disk Group

If you want to use ASM as the storage option for either database or recovery files, and an existing ASM disk group exists, you have the following choices, depending on the installation method that you select:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to create a new disk group or use an existing one.
The same choice is available to you if you use DBCA after the installation to create a database.
- If you select an installation type that runs DBCA in non-interactive mode, you must choose an existing disk group for the new database; you cannot create a new disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The ASM instance that manages the existing disk group can be running in a different Oracle home directory.

To determine whether an existing ASM disk group exists, or to determine whether there is sufficient disk space in a disk group, you can use Oracle Enterprise Manager Database Grid Control or Database Control. Alternatively, you can use the following procedure:

1. View the contents of the `oratab` file to determine whether an ASM instance is configured on the system:

```
# more /etc/oratab
```

If an ASM instance is configured on the system, the `oratab` file should contain a line similar to the following:

```
+ASM:oracle_home_path:N
```

In this example, `+ASM` is the system identifier (SID) of the ASM instance and `oracle_home_path` is the Oracle home directory where it is installed. By convention, the SID for an ASM instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the ASM instance that you want to use.
3. Connect to the ASM instance as the `SYS` user with `SYSDBA` privilege and start the instance if necessary:

```
# $ORACLE_HOME/bin/sqlplus "SYS/SYS_password as SYSDBA"  
SQL> STARTUP
```

4. Enter the following command to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

```
SQL> SELECT NAME,TYPE,TOTAL_MB,FREE_MB FROM V$ASM_DISKGROUP;
```

5. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.
6. If necessary, install or identify the additional disk devices required to meet the storage requirements listed in the previous section.

Note: If you are adding devices to an existing disk group, Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Configuring Disks for ASM

To configure disks for use with ASM, follow these steps:

1. If necessary, install the shared disks that you intend to use for the ASM disk group.
2. To make sure that the disks are available, enter the following command:

```
# /usr/sbin/ioscan -fun -C disk
```

The output from this command is similar to the following:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	0	0/0/1/0.6.0	sdisk	CLAIMED	DEVICE	HP DVD-ROM 6x/32x /dev/dsk/c0t6d0 /dev/rdisk/c0t6d0
disk	1	0/0/1/1.2.0	sdisk	CLAIMED	DEVICE	SEAGATE ST39103LC /dev/dsk/c1t2d0 /dev/rdisk/c1t2d0

This command displays information about each disk attached to the system, including the block device name (`/dev/dsk/cxydz`) and the character raw device name (`/dev/rdisk/cxydz`).

If this command does not display device name information, enter the following command to install the special device files:

```
# /usr/sbin/insf -e
```

3. For each disk that you want to add to a disk group, enter the following command to verify that it is not already part of an LVM volume group:

```
# /sbin/pvdisplay /dev/dsk/cxydz
```

If this command displays volume group information, the disk is already part of a volume group. The disks that you choose must not be part of an LVM volume group.

Note: If you are using different volume management software, for example VERITAS Volume Manager, see the appropriate documentation for information about verifying that a disk is not in use.

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

```
# chown oracle:dba /dev/rdisk/cxydz
# chmod 660 /dev/rdisk/cxydz
```

Note: If you are using a multi-pathing disk driver with ASM, ensure that you set the permissions only on the correct logical device name for the disk.

Configure Raw Logical Volumes

See the following sections for information about configuring raw logical volumes:

- ["Review Important Information"](#) on page 3-28
- ["Disable Operating System Activation of Shared Volume Groups"](#) on page 3-29
- ["Create a Volume Group"](#) on page 3-30
- ["Create Raw Logical Volumes in the New Volume Group"](#) on page 3-31
- ["Export the Volume Group and Import It on the Other Cluster Nodes"](#) on page 3-33
- ["Activate the Volume Group in Shared Mode on All Cluster Nodes"](#) on page 3-34
- ["Create the DBCA Raw Device Mapping File"](#) on page 3-34

Note: If you are using ASM for database file storage, you need only create raw devices for the Oracle CRS files.

The following subsections describe how to create logical volumes using the command line. If you prefer, you can use SAM to complete the same tasks. See the HP-UX documentation for more information about using SAM.

Review Important Information

This section describes how to configure shared raw logical volumes for Oracle CRS and database file storage. The procedures in this section describe how to create a new shared volume group that contains the logical volumes required for both types of files.

To use shared raw logical volumes, HP Serviceguard Extensions for RAC must be installed on all cluster nodes. If HP Serviceguard Extensions for RAC is not installed, you can use raw partitions on a shared disk to store the Oracle CRS files. However, Oracle recommends that you use this method only for the CRS files and use an alternative method such as ASM for database file storage.

Before you continue, review the following guidelines which contain important information about using shared logical volumes with this release of Oracle Real Application Clusters:

- You must use shared volume groups for Oracle CRS and database files.
- If you specify raw logical volumes for the Oracle Cluster Registry and Oracle CRS voting disk during the Oracle CRS installation, the installation process configures the Oracle CRS startup scripts to activate the volume group that contains these logical volumes when Oracle CRS starts.

Note: You do not need to create a Serviceguard package to activate the volume group that contains the Oracle CRS files.

- The Oracle CRS files require less than 200 MB of disk space. To make efficient use of the disk space in a volume group, Oracle recommends that you use the same shared volume group for the logical volumes for both the CRS files and the database files.
- If you are upgrading an existing Oracle9*i* release 2 RAC installation that uses raw logical volumes, you can use the existing SRVM configuration repository logical volume for the OCR and create a new logical volume in the same volume group for the Oracle CRS voting disk. However, before you install Oracle CRS, you must

remove this volume group from any Serviceguard package that currently activates it.

See Also: For information about removing a volume group from a Serviceguard package, see the HP Serviceguard or HP Serviceguard Extensions for RAC documentation.

Note: If you are upgrading a database, you must also create a new logical volume for the SYSAUX tablespace. See the "[Create Raw Logical Volumes in the New Volume Group](#)" section on page 3-31 for more information about the requirements for the Oracle CRS voting disk and SYSAUX logical volumes.

- You must use either your own startup script or a Serviceguard package to activate new or existing volume groups that contain only database files (not Oracle CRS files).

See Also: For information about creating a Serviceguard package to activate a shared volume group for RAC, see the HP Serviceguard documentation.

- All shared volume groups that you intend to use for Oracle CRS or database files must be activated in shared mode before you start the installation.
- All shared volume groups that you are using for RAC, including the volume group that contains the Oracle CRS files, must be specified in the cluster configuration file using the parameter OPS_VOLUME_GROUP.

Note: If you create a new shared volume group for RAC on an existing HP Serviceguard cluster, you must reconfigure and restart the cluster before installing Oracle CRS. See the HP Serviceguard documentation for information about configuring the cluster and specifying shared volume groups.

- The procedures in this section describe how to create basic volumes groups and volumes. If you want to configure more complex volumes, using mirroring for example, use this section in conjunction with the HP Serviceguard documentation.

Disable Operating System Activation of Shared Volume Groups

To prevent the operating system from activating shared volume groups when it boots, you must edit the `/etc/lvmrc` file, as follows:

1. Create a backup copy of the `/etc/lvmrc` file:

```
# cp /etc/lvmrc /etc/lvmrc_orig
```

2. Open the `/etc/lvmrc` file in any text editor and search for the `AUTO_VG_ACTIVATE` flag.
3. If necessary, change the value of the `AUTO_VG_ACTIVATE` flag to 0, to disable automatic volume group activation, as follows:

```
AUTO_VG_ACTIVATE=0
```

4. Search for the `custom_vg_activation` function in the `/etc/lvmrc` file.
5. Add `vgchange` commands to the function, as shown in the following example, to automatically activate existing local volume groups:

```
custom_vg_activation()
{
    # e.g. /sbin/vgchange -a y -s
    #     parallel_vg_sync "/dev/vg00 /dev/vg01"
    #     parallel_vg_sync "/dev/vg02 /dev/vg03"

    /sbin/vgchange -a y vg00
    /sbin/vgchange -a y vg01
    /sbin/vgchange -a y vg02

    return 0
}
```

In this example, `vg00`, `vg01`, and `vg02` are the volume groups that you want to activate automatically when the system reboots.

Note: You must not specify any shared volume groups in this file. The shared volume group that you use for Oracle CRS files will be activated by the CRS startup scripts.

Create a Volume Group

To create a volume group, follow these steps:

1. If necessary, install the shared disks that you intend to use for the database.
2. To make sure that the disks are available, enter the following command:

```
# /sbin/ioscan -fun -C disk
```

The output from this command is similar to the following:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	0	0/0/1/0.6.0	sdisk	CLAIMED	DEVICE	HP DVD-ROM 6x/32x
					/dev/dsk/c0t6d0	/dev/rdisk/c0t6d0
disk	1	0/0/1/1.2.0	sdisk	CLAIMED	DEVICE	SEAGATE ST39103LC
					/dev/dsk/c1t2d0	/dev/rdisk/c1t2d0
disk	2	0/0/2/0.2.0	sdisk	CLAIMED	DEVICE	SEAGATE ST118202LC
					/dev/dsk/c2t2d0	/dev/rdisk/c2t2d0

This command displays information about each disk attached to the system, including the block device name (`/dev/dsk/cxydz`) and the character raw device name (`/dev/rdisk/cxydz`).

If this command does not display device name information for a disk, enter the following command to install the special device files:

```
# /usr/sbin/insf -e
```

3. For each disk that you want to add to the volume group, enter the following command to verify that it is not already part of a volume group:

```
# /sbin/pvdisplay /dev/dsk/cxydz
```

If this command displays volume group information, the disk is already part of a volume group.

4. For each disk that you want to add to the volume group, enter a command similar to the following:

```
# /sbin/pvcreate /dev/rdisk/cxytdz
```

5. To create a directory for the volume group in the `/dev` directory, enter a command similar to the following, where `vg_name` is the name you want to use for the volume group:

```
# mkdir /dev/vg_name
```

6. To identify used device minor numbers, enter the following command on each node of the cluster:

```
# ls -la /dev/*/group
```

This command displays information about the device numbers used by all configured volume groups, similar to the following:

```
crw-r----- 1 root  sys      64 0x000000 Mar  4 2002 /dev/vg00/group
crw-r--r--  1 root  sys      64 0x010000 Mar  4 2002 /dev/vg01/group
```

In this example, 64 is the major number used by all volume group devices and 0x000000 and 0x010000 are the minor numbers used by volume groups `vg00` and `vg01` respectively. Minor numbers have the format `0xnn0000`, where `nn` is a number in the range 00 to the value of the `maxvgs` kernel parameter minus 1. The default value for the `maxvgs` parameter is 10, so the default range is 00 to 09.

7. Identify an appropriate minor number that is unused on all nodes in the cluster.
8. To create the volume group and activate it, enter commands similar to the following:

```
# /sbin/mknod /dev/vg_name/group c 64 0xnn0000
# /sbin/vgcreate /dev/vg_name /dev/dsk/cxytdz . . .
# /sbin/vgchange -a y vg_name
```

In this example:

- `vg_name` is the name you want to give the volume group
- `0xnn0000` is a minor number that is unused on all nodes in the cluster.
- `/dev/dsk/cxytdz . . .` is a list of one or more block device names for the disks that you want to add to the volume group.

Create Raw Logical Volumes in the New Volume Group

To create the required raw logical volumes in the new volume group, follow these steps:

1. Choose a name for the database that you want to create.

The name that you choose must start with a letter and have no more than four characters, for example, `orcl`.

2. From the following table, identify the logical volumes that you must create:

Number	Partition Size (MB)	Purpose and Sample Logical Volume Name
Raw Logical Volumes for Oracle Database		
1	500	SYSTEM tablespace: <i>dbname_system_raw_500m</i>
1	300 + (Number of instances * 250)	SYSAUX tablespace: <i>dbname_sysaux_raw_800m</i>
Number of instances	500	UNDOTBS n tablespace (One tablespace for each instance, where n is the number of the instance): <i>dbname_undotbsn_raw_500m</i>
1	160	EXAMPLE tablespace: <i>dbname_example_raw_160m</i>
1	120	USERS tablespace: <i>dbname_users_raw_120m</i>
2 * number of instances	120	Two online redo log files for each instance (where n is the number of the instance and m is the log number, 1 or 2): <i>dbname_redonm_raw_120m</i>
2	110	First and second control files: <i>dbname_control[1 2]_raw_110m</i>
1	250	TEMP tablespace: <i>dbname_temp_raw_250m</i>
1	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwdfile_raw_5m</i>
Raw Logical Volumes for Oracle Cluster Ready Services (CRS)		
1	100	Oracle Cluster Registry: <i>ora_ocr_raw_100m</i> Note: You need to create this raw logical volume only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle cluster registry. If you are upgrading from Oracle9i Release 2, you can continue to use the raw device that you used for the SRVM configuration repository instead of creating this new logical volume.
1	20	Oracle CRS voting disk: <i>ora_vote_raw_20m</i> Note: You need to create this raw logical volume only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle CRS voting disk.

- To create each required logical volume, enter a command similar to the following:

```
# /sbin/lvcreate -n LVname -L size /dev/VGname
```

In this example:

- *LVname* is the name of the logical volume that you want to create
Oracle recommends that you use the sample names shown in the previous table for the logical volumes. Substitute the *dbname* variable in the sample logical volume name with the name you choose for the database in step 1.
- *VGname* is the name of the volume group where you want to create the logical volume
- *size* is the size of the logical volume in megabytes

The following example shows a sample command used to create an 800 MB logical volume in the `oracle_vg` volume group for the SYSAUX tablespace of a database named `test`:

```
# /sbin/lvcreate -n test_sysaux_800m -L 800 /dev/oracle_vg
```

4. Change the owner, group, and permissions on the character device files associated with the logical volumes that you created, as follows:

```
# chown oracle:dba /dev/vg_name/r*
# chmod 755 /dev/vg_name
# chmod 660 /dev/vg_name/r*
```

5. Change the owner and group on the character device file associated with the logical volume for the Oracle Cluster Registry, as follows:

```
# chown root:dba /dev/vg_name/rocr_ocr_raw_100m
```

Export the Volume Group and Import It on the Other Cluster Nodes

To export the volume group and import it on the other cluster nodes, follow these steps:

1. Deactivate the volume group:

```
# /sbin/vgchange -a n vg_name
```

2. To export the description of the volume group and its associated logical volumes to a map file, enter a command similar to the following:

```
# /sbin/vgexport -v -s -p -m /tmp/vg_name.map /dev/vg_name
```

3. Enter commands similar to the following to copy the map file to the other cluster nodes:

```
# rcp /tmp/vg_name.map nodename:/tmp/vg_name.map
```

4. Enter commands similar to the following on the other cluster nodes to import the volume group that you created on the first node:

```
# mkdir /dev/vg_name
# /sbin/mknod /dev/vg_name/group c 64 0xnn0000
# /sbin/vgimport -v -s -m /tmp/vg_name.map /dev/vg_name
```

5. Enter commands similar to the following on the other cluster nodes to change the owner, group, and permissions on the character device files associated with the logical volumes that you created:

```
# chown oracle:dba /dev/vg_name/r*
# chmod 755 /dev/vg_name
# chmod 660 /dev/vg_name/r*
```

6. Change the owner and group on the character device file associated with the logical volume for the Oracle Cluster Registry, as follows:

```
# chown root:dba /dev/vg_name/rora_ocr_raw_100m
```

Activate the Volume Group in Shared Mode on All Cluster Nodes

To activate the volume group in shared mode on all cluster nodes, enter the following command on each node:

```
# /sbin/vgchange -a s vg_name
```

Note: After you install Oracle CRS, the CRS startup scripts use this command to activate the volume group automatically when the node boots.

Create the DBCA Raw Device Mapping File

Note: You must complete this procedure only if you are using raw logical volumes for database files. You do not specify the raw logical volumes (or partitions) for the Oracle CRS files in the DBCA raw device mapping file.

To enable Database Configuration Assistant (DBCA) to identify the appropriate raw device for each database file, you must create a raw device mapping file, as follows:

1. Set the ORACLE_BASE environment variable to specify the Oracle base directory that you identified or created previously:

- Bourne, Bash or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle ; export ORACLE_BASE
```

- C shell:

```
$ setenv ORACLE_BASE /u01/app/oracle
```

2. Create a database file subdirectory under the Oracle base directory and set the appropriate owner, group, and permissions on it:

```
# mkdir -p $ORACLE_BASE/oradata/dbname
# chown -R oracle:oinstall $ORACLE_BASE/oradata
# chmod -R 775 $ORACLE_BASE/oradata
```

In this example, *dbname* is the name of the database that you chose previously.

3. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
4. Enter a command similar to the following to create a text file that you can use to create the raw device mapping file:

```
# find /dev/vg_name -user oracle -name 'r*' -print > dbname_raw.conf
```

5. Edit the `dbname_raw.conf` file in any text editor to create a file similar to the following:

Note: The following example shows a sample mapping file for a two-instance RAC cluster.

```

system=/dev/vg_name/rdbname_system_raw_500m
sysaux=/dev/vg_name/rdbname_sysaux_raw_800m
example=/dev/vg_name/rdbname_example_raw_160m
users=/dev/vg_name/rdbname_users_raw_120m
temp=/dev/vg_name/rdbname_temp_raw_250m
undotbs1=/dev/vg_name/rdbname_undotbs1_raw_500m
undotbs2=/dev/vg_name/rdbname_undotbs2_raw_500m
redo1_1=/dev/vg_name/rdbname_redo1_1_raw_120m
redo1_2=/dev/vg_name/rdbname_redo1_2_raw_120m
redo2_1=/dev/vg_name/rdbname_redo2_1_raw_120m
redo2_2=/dev/vg_name/rdbname_redo2_2_raw_120m
control1=/dev/vg_name/rdbname_control1_raw_110m
control2=/dev/vg_name/rdbname_control2_raw_110m
spfile=/dev/vg_name/rdbname_spfile_raw_5m
pwdfile=/dev/vg_name/rdbname_pwdfile_raw_5m

```

In this example:

- *vg_name* is the name of the volume group
- *dbname* is the name of the database

Use the following guidelines when creating or editing this file:

- Each line in the file must have the following format:

```
database_object_identifier=logical_volume
```

The logical volume names suggested in this manual include the database object identifier that you must use in this mapping file. For example, in the following logical volume name, *redo1_1* is the database object identifier:

```
/dev/oracle_vg/rac_redo1_1_raw_120m
```

- For a single instance database, the file must specify one automatic undo tablespace data file (*undotbs1*), and at least two redo log files (*redo1_1*, *redo1_2*).
 - For a RAC database, the file must specify one automatic undo tablespace datafile (*undotbsn*) and two redo log files (*redon_1*, *redon_2*) for each instance.
 - Specify at least two control files (*control1*, *control2*).
 - To use manual instead of automatic undo management, specify a single RBS tablespace datafile (*rbs*) instead of the automatic undo management tablespaces.
6. Save the file and note the file name that you specified.
 7. When you are configuring the `oracle` user's environment later in this chapter, set the `DBCA_RAW_CONFIG` environment variable to specify the full path to this file.

Verify the Cluster Software Configuration

Note: The cluster software mentioned in the following sections is not required to install Oracle Real Application Clusters. However, if this software is installed, Oracle Cluster Ready Services (CRS) can integrate with it.

To verify that the HP Serviceguard software is running:

1. If necessary, configure (or reconfigure) the cluster.
2. Verify that HP Serviceguard is configured to start automatically at runlevel 3.
3. To verify that HP Serviceguard is running correctly, enter the following command:

```
# /usr/sbin/cmviewcl -v
```

Verify that the cluster status is up, and that all nodes are up and running.

4. If the cluster is not running, enter the following command to start it:
5. If any nodes are not running, enter a command similar to the following to start the cluster daemon on those nodes:

```
# /usr/sbin/cmrunnode nodename1 nodename2...
```

If you intend to use Hyper Messaging protocol as the cluster interconnect, enter the following command to verify that the HyperFabric software is configured properly:

```
# /opt/clic/bin/clic_stat -d NET
```

In the output from this command, verify that all of the cluster nodes are listed in the Node Table Entries section. If a node is missing, see the HyperFabric documentation for more information about configuring the software.

Stop Existing Oracle Processes

Caution: If you are installing additional Oracle Database 10g products in an existing Oracle home, stop **all** processes running in the Oracle home. You must complete this task to enable the Installer to relink certain executables and libraries.

If you choose to create a database during the installation, most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the IPC key value EXTPROC. However, if an existing Oracle Net listener process is using the same port or key value, the Installer can only configure the new listener; it cannot start it. To ensure that the new listener process starts during the installation, you must shut down any existing listeners before starting the Installer.

To determine whether an existing listener process is running and to shut it down if necessary, follow these steps:

1. Switch user to oracle:

```
# su - oracle
```

2. Enter the following command to determine whether a listener process is running and to identify its name and the Oracle home directory in which it is installed:

```
$ ps -ef | grep tnslnsr
```

This command displays information about the Oracle Net listeners running on the system:

```
... oracle_home1/bin/tnslnsr LISTENER -inherit
```

In this example, *oracle_home1* is the Oracle home directory where the listener is installed and LISTENER is the listener name.

3. Set the ORACLE_HOME environment variable to specify the appropriate Oracle home directory for the listener:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=oracle_home1
$ export ORACLE_HOME
```

- C or tcsh shell:

```
$ setenv ORACLE_HOME oracle_home1
```

4. Enter the following command to identify the TCP/IP port number and IPC key value that the listener is using:

```
$ $ORACLE_HOME/bin/lsnrctl status listenername
```

Note: If the listener uses the default name LISTENER, you do not have to specify the listener name in this command.

5. Enter a command similar to the following to stop the listener process:

```
$ $ORACLE_HOME/bin/lsnrctl stop listenername
```

6. Repeat this procedure to stop all listeners running on this system and on all other cluster nodes.

Configure the oracle User's Environment

You run the Installer from the `oracle` account. However, before you start the Installer you must configure the environment of the `oracle` user. To configure the environment, you must:

- Set the default file mode creation mask (`umask`) to 022 in the shell startup file.
- Set the `DISPLAY` and `ORACLE_BASE` environment variables.

To set the `oracle` user's environment, follow these steps:

1. Start a new terminal session, for example, an X terminal (`xterm`).
2. Enter the following command to ensure that X Window applications can display on this system:

```
$ xhost +
```

3. If you are not already logged in to the system where you want to install the software, log in to that system as the `oracle` user.

4. If you are not logged in as the `oracle` user, switch user to `oracle`:

```
$ su - oracle
```
5. To determine the default shell for the `oracle` user, enter the following command:

```
$ echo $SHELL
```
6. Open the `oracle` user's shell startup file in any text editor:
 - Bourne shell (`sh`), Bash shell (`bash`), or Korn shell (`ksh`):

```
$ vi .profile
```
 - C shell (`csh` or `tcsh`):

```
% vi .login
```
7. Enter or edit the following line, specifying a value of 022 for the default file creation mask:

```
umask 022
```
8. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, remove the appropriate lines from the file.
9. Save the file and exit from the editor.
10. To run shell startup script, enter one of the following command:
 - Bourne, Bash, or Korn shell:

```
$ . ~/.profile
```
 - C shell:

```
% source ~/.login
```
11. If you are not installing the software on the local system, enter a command similar to the following to direct X applications to display on the local system:
 - Bourne, Bash, or Korn shell:

```
$ DISPLAY=local_host:0.0 ; export DISPLAY
```
 - C shell:

```
$ setenv DISPLAY local_host:0.0
```

In this example, `local_host` is the host name or IP address of the system you want to use to display the Installer (your workstation or PC).
12. If you determined that the `/tmp` directory has less than 200 MB of free disk space, identify a file system with at least 200 MB of free space and set the `TEMP` and `TMPDIR` environment variables to specify a temporary directory on this file system:
 - a. Use the `bd` command to identify a suitable file system with sufficient free space.
 - b. If necessary, enter commands similar to the following to create a temporary directory on the file system you identified, and set the appropriate permissions on the directory:

```
$ su - root  
# mkdir /mount_point/tmp
```

```
# chmod a+wr /mount_point/tmp
# exit
```

- c. Enter commands similar to the following to set the TEMP and TMPDIR environment variables:

- * Bourne, Bash, or Korn shell:

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

- * C shell:

```
$ setenv TEMP /mount_point/tmp
$ setenv TMPDIR /mount_point/tmp
```

13. Enter commands similar to the following to set the ORACLE_BASE environment variable:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle
$ export ORACLE_BASE
```

- C shell:

```
% setenv ORACLE_BASE /u01/app/oracle
```

In these examples, /u01/app/oracle is the Oracle base directory that you created earlier.

14. If necessary, set the following environment variables. If you must specify several values for an environment variable, for example PATH, separate the values with a colon (:).

Environment Variable	Required By	Sample Setting and Description
DBCA_RAW_CONFIG	Oracle Database installations that use raw devices for database storage.	\$ORACLE_BASE/oradata/dbname/dbname_raw.conf Specifies the location of the raw device mapping file
COBDIR	Pro*COBOL	/opt/lib/cobol Specifies the directory where COBOL is installed on the system. Note: To install Pro*COBOL, you must choose a custom software installation.
PATH	Pro*COBOL	\$COBDIR/bin:\$PATH Specifies the directory where the COBOL compiler executable is located.

15. Enter the following command to ensure that the ORACLE_HOME and TNS_ADMIN environment variable are not set:

```
$ unset ORACLE_HOME
$ unset TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value it specifies as the default path for the Oracle home directory. However, if you set the ORACLE_BASE environment variable, Oracle recommends that you unset this environment variable and choose the default path suggested by the Installer.

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

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

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

Pre-Installation Tasks for Installing RAC on HP Tru64 UNIX

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer to install RAC on HP Tru64 UNIX systems. This chapter includes following topics:

- Log In to the System as root
- Check the Hardware Requirements
- Check the Network Requirements
- Check the Software Requirements
- Create Required UNIX Groups and User
- Configure Kernel Subsystem Attributes
- Identify Required Software Directories
- Identify or Create an Oracle Base Directory
- Configure Disk Storage for Oracle CRS and Database Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Verify that the Cluster Software is Running
- Stop Existing Oracle Processes
- Configure the oracle User's Environment

Log In to the System as root

Before you install the Oracle software, you must complete several tasks as the `root` user. To log in as the `root` user, complete one of the following procedures:

Note: Unless you intend to complete a silent installation, you must install the software from an X Window System workstation, an X terminal, or a PC or other system with X server software installed.

For more information about non-interactive and silent installations, see the appendices in *Oracle Database Installation Guide for UNIX Systems*.

- If you are installing the software from an X Window System workstation or X terminal:
 1. Start a local terminal session, for example, an X terminal (`xterm`).
 2. If you are not installing the software on the local system, enter the following command to enable remote hosts to display X applications on the local X server:

```
$ xhost +
```
 3. If you want to install the software on a remote system, enter a command similar to the following to connect to that system:

```
$ telnet remote_host
```
 4. If you are not logged in as the `root` user, enter the following command to switch user to `root`:

```
$ su - root
password:
#
```
- If you are installing the software from a PC or other system with X server software installed:

Note: If necessary, see your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

1. Start the X server software.
2. Configure the security settings of the X server software to permit remote hosts to display X applications on the local system.
3. Connect to the remote system where you want to install the software and start a terminal session on that system, for example, an X terminal (`xterm`).
4. If you are not logged in as the `root` user on the remote system, enter the following command to switch user to `root`:

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

Check the Hardware Requirements

The system must meet the following minimum hardware requirements:

- 512 MB of physical RAM
- 1 GB of swap space (or twice the size of RAM)
On systems with 2 GB or more of RAM, the swap space can be between one and two times the size of RAM.
- 400 MB of disk space in the `/tmp` directory
- Up to 4 GB of disk space for the Oracle software, depending on the installation type
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirements for databases that use Automatic Storage Management (ASM) or raw device storage is described later in this chapter.

Additional disk space, either on a file system or in an ASM disk group, is required for the flash recovery area if you choose to configure automated backups.

To ensure that the system meets these requirements, follow these steps:

1. To determine the physical RAM size, enter the following command:

```
# /bin/vmstat -P | grep "Total Physical Memory"
```

 If the size of the physical RAM installed in the system is less than the required size, you must install more memory before continuing.
2. To determine the size of the configured swap space, enter the following command:

```
# /sbin/swapon -s
```

 If necessary, see your operating system documentation for information about how to configure additional swap space.
3. To determine the amount of disk space available in the `/tmp` directory, enter the following command:

```
# df -k /tmp
```

If there is less than 400 MB of disk space available in the `/tmp` directory, complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to achieve the required disk space.
 - Set the `TEMP` and `TMPDIR` environment variables when setting the `oracle` user's environment (described later).
 - Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.
4. To determine the amount of free disk space on the system, enter the following command:

df -k

Check the Network Requirements

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Real Application Clusters installation.

Note: For the most up-to-date information about supported network protocols and hardware for RAC installations, see the Certify pages on the Oracle *MetaLink* Web site:

<http://metalink.oracle.com>

Network Hardware Requirements

Each node in the cluster must meet the following requirements:

- Each node must have at least two network adapters; one for the public network interface and one for the private network interface (the interconnect).
- The interface names associated with the network adapters for each network must be the same on all nodes.
- For increased reliability, you can configure redundant public and private network adapters for each node.
- For the public network, each network adapter must support TCP/IP.
- For the private network, Oracle supports the following interconnect protocols and hardware:
 - Reliable datagram (RDG) using Memory Channel adapters and hubs
 - RDG or user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended)

Note: RDG is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS. To use UDP as the interconnect protocol for RAC, you must relink the `oracle` executable. See the *Oracle Database Administrator's Reference for UNIX Systems* for more information about enabling UDP for RAC on Tru64 UNIX systems.

IP Address Requirements

Before starting the installation, you must identify or obtain the following IP addresses for each node:

- An IP address and an associated host name registered in the domain name service (DNS) for each public network interface
- One unused virtual IP address and an associated virtual host name registered in DNS that you will configure for the primary public network interface

The virtual IP address must be in the same subnet as the associated public interface. After installation, you can configure clients to use the virtual host name or IP address. If a node fails, its virtual IP address fails over to another node.

- A private IP address and optional host name for each private interface

Oracle recommends that you use private network IP addresses for these interfaces, for example: 10.*.* or 192.168.*.*. You can use the `/etc/hosts` file on each node to associate private host names with private IP addresses.

For example, if each node has two public and two private interfaces, you might have the following host names and IP addresses on one of the nodes (`rac1`) and similar host names and IP addresses on the other nodes:

Host Name	Type	IP Address	Registered In
<code>rac1.mydomain.com</code>	Public	143.47.43.100	DNS
<code>rac1-2.mydomain.com</code>	Public	143.46.51.101	DNS
<code>rac1-vip.mydomain.com</code>	Virtual	143.46.43.104	DNS
<code>rac1-priv1</code>	Private	10.0.0.1	<code>/etc/hosts</code>
<code>rac1-priv2</code>	Private	10.0.0.2	<code>/etc/hosts</code>

Checking the Network Requirements

To verify that each node meets the requirements, follow these steps:

1. If necessary, install the network adapters for the public and private networks and configure them with either public or private IP addresses.
2. Register the host names and IP addresses for the public network interfaces in DNS.
3. For each node, register one virtual host name and IP address in DNS.
4. For each private interface on every node, add a line similar to the following to the `/etc/hosts` file on all nodes, specifying the private IP address and associated private host name:

```
10.0.0.1    rac1-priv1
```

5. To identify the interface name and associated IP address for every network adapter, enter the following command:

```
# /sbin/ifconfig -a
```

From the output, identify the interface name and IP address for all network adapters that you want to specify as public or private network interfaces.

Note: When you install Oracle CRS and RAC, you will require this information.

Check the Software Requirements

Note: The Oracle Universal Installer performs checks on your system to verify that it meets the requirements listed for HP Tru64. To ensure that these checks pass, verify the requirements before you start the Installer.

Depending on the products that you intend to install, verify that the following software is installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	Operating system version: HP Tru64 UNIX V5.1B <hr/> Software Development Kit (SDK) v 1.4.2 for the Tru64 UNIX Operating System for the Java Platform (JDK 1.4.2) <hr/> Operating system subsets: OSFCMPLRS OSFLIBA OSFPGMR OSFSER OSFX11
Oracle Real Application Clusters	HP TruCluster 5.1B: TCRBASE TCRMIGRATE
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE	DCE v4.1 or later: DCECDS DECRTS DCESEC
Oracle Spatial	X Window and X/Motif software: OSFXDEVX OSFXINCLUDE OSFXLIBA Note: This software is required only to build the sample programs.
Oracle Messaging Gateways	IBM MQSeries V5.1 client and server: MQS_CLIENT MQS_SERVER MQSeries classes for Java and MQSeries classes for Java Message Service V5.2.2 (SupportPac MA88): MQS_MQJAVA
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Compaq C Compiler V6.5-207 (dtk): DTCCMPLR Compaq C++ Version 6.5-014: CXXBASE CXXLIB CXXOLD
Oracle JDBC/OCI Drivers	You can use the following optional JDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation: <ul style="list-style-type: none"> ■ SDK v 1.3.1-5 ■ SDK v 1.2.2-12 Note: SDK v 1.4.2 must be installed before you install the software.
Oracle Transparent Gateway for Sybase	Either of the following: <ul style="list-style-type: none"> ■ Sybase Adaptive Server Enterprise V12 or 12.5 ■ Sybase Open Client V12 or 12.5 Open Client is required if the Sybase server is not on the local node.

To ensure that the system meets these requirements, follow these steps:

1. To determine which version of Tru64 UNIX is installed, enter the following command:

```
# /usr/sbin/sizer -v
Compaq Tru64 UNIX V5.1B (Rev. 2650); Mon Nov  3 10:13:28 PST 200
```

In this example, the version shown is V5.1B. If necessary, see your operating system documentation for information about upgrading the operating system.

2. To determine whether Java SDK 1.4.2 is installed, enter the following command:

```
# /usr/sbin/setld -i JAVA142 | more
```

If Java SDK 1.4.2 is installed, this command displays the paths to all of the installed files. Note the path of the Java home directory. You must specify this value during the installation. The default path is:

```
/usr/opt/java142
```

If this command returns the message Unknown subset, Java SDK 1.4.2 is not installed. Download the Java SDK 1.4.2.01 or higher from the following Web site and install it:

```
http://www.compaq.com/java/download/index.html
```

3. To determine whether the required software subsets are installed, enter one of the following commands:

- To view the list of all software subsets installed on the system, enter the following command:

```
# /usr/sbin/setld -i | more
```

- To determine whether a particular software subset is installed, enter a command similar to the following:

```
# /usr/sbin/setld -i | grep subsetname
```

If necessary, install the required software subset. If you require the Compaq C Compiler V6.5-207 (dtk), you can download it from the following Web site:

```
http://www.tru64unix.compaq.com/dtk/
```

4. If you intend to use Oracle Messaging Gateway and require MQSeries classes for Java and MQSeries classes for Java Message Service (SupportPac MA88), download it from the following Web site:

```
http://www.ibm.com/software/integration/support/supportpacs/individual/ma88.html
```

Check for Required Patches

Depending on the products that you intend to install, verify that the following patches are installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	Tru64 UNIX V5.1B Patch Kit 2 or higher: T64V51BB22AS0002-20030415 HP Tru64 UNIX 5.1B PK2 BL22 Fixes for AdvFS Panic in _OtsMove; and Possible Memory Corruption: T64KIT0020879-V51BB22-E-20031125
Oracle Messaging Gateway	Corrective service diskette (CSD) for MQSeries: CSD09 or later for MQSeries V5.1

To determine whether the required patch kits are installed, enter the following command:

```
# /usr/sbin/dupatch -track -type kit
```

If this command does not display the identifiers shown in the previous table for the required patch kits (or the identifier for a higher patch kit level), download the latest patch kit from the following Web site and install it (registration is required to access this Web site):

<http://itrc.hp.com/service/patch/mainPage.do>

If you require a CSD for MQSeries, see the following Web site for download and installation information:

<http://www.ibm.com/software/integration/mqfamily/support/summary/dig.html>

Create Required UNIX Groups and User

Depending on whether this is the first time Oracle software is being installed on this system and on the products that you are installing, you may need to create several UNIX groups and a UNIX user.

The following UNIX groups and user are required if you are installing Oracle Database:

- The OSDBA group (`dba`)

You must create this group the first time you install Oracle Database software on the system. It identifies UNIX users that have database administrative privileges (the SYSDBA privilege). The default name for this group is `dba`.

If you want to specify a group name other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of this group. In this case, the Installer prompts you to specify the name of this group.

- The OSOPER group (`oper`)

This is an optional group. Create this group if you want a separate group of UNIX users to have a limited set of database administrative privileges (the SYSOPER privilege). By default, members of the OSDBA group also have the SYSOPER privilege.

If you want to specify a separate OSOPER group, other than the default `dba` group, you must choose the Custom installation type to install the software or

start the Installer as a user that is not a member of the `dba` group. In this case, the Installer prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

- An unprivileged user (`nobody`)

You must verify that the unprivileged user `nobody` exists on the system. The `nobody` user must own the external jobs (`ext job`) executable after the installation.

The following UNIX group and user are required for all installations:

- The Oracle Inventory group (`oinstall`)

You must create this group the first time you install Oracle software on the system. The usual name chosen for this group is `oinstall`. This group owns the Oracle inventory, which is a catalog of all Oracle software installed on the system.

Note: If Oracle software is already installed on the system, the existing Oracle Inventory group must be the primary group of the UNIX user that you use to install new Oracle software. The following sections describe how to identify an existing Oracle Inventory group.

- The Oracle software owner user (`oracle`)

You must create this user the first time you install Oracle software on the system. This user owns all of the software installed during the installation. The usual name chosen for this user is `oracle`. This user must have the Oracle Inventory group as its primary group. It must also have the OSDBA and OSOPER groups as a secondary groups.

A single Oracle Inventory group is required for all installations of Oracle software on the system. After the first installation of Oracle software, you must use the same Oracle Inventory group for all subsequent Oracle software installations on that system. However, you can choose to create different Oracle software owner users, OSDBA groups, and OSOPER groups (other than `oracle`, `dba`, and `oper`) for separate installations. By using different groups for different installations, members of these different groups have DBA privileges only on the associated databases rather than on all databases on the system.

See Also: The *Oracle Database Administrator's Reference for UNIX Systems* and the *Oracle Database Administrator's Guide* contain more information about the OSDBA and OSOPER groups and the SYSDBA and SYSOPER privileges.

The following sections describe how to create the required UNIX user and groups.

Note: The following sections describe how to create local users and groups. As an alternative to creating local users and groups, you could create the appropriate users and groups in a directory service, for example, Network Information Services (NIS). For information about using directory services, contact your system administrator or see your operating system documentation.

Creating the Oracle Inventory Group

You must create the Oracle Inventory group if it does not already exist. The following subsections describe how to determine the Oracle Inventory group name, if it exists, and how to create it if necessary.

Determine Whether the Oracle Inventory Group Exists

When you install Oracle software on the system for the first time, the Installer creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group and the path of the Oracle Inventory directory. To determine whether the Oracle Inventory group exists, enter the following command:

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

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inst_group` parameter shows the name of the Oracle Inventory group (`oinstall`).

Create the Oracle Inventory Group

If the `oraInst.loc` file does not exist, create the Oracle Inventory group, `oinstall`, by entering the following command:

```
# /usr/sbin/groupadd oinstall
```

Creating the OSDBA Group

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system
- An OSBDA group exists, but you want to give a different group of UNIX users database administrative privileges in a new Oracle installation

If the OSDBA group does not exist or if you require a new OSDBA group, create it by entering the following command, using the group name `dba` unless a group with that name already exists:

```
# /usr/sbin/groupadd dba
```

Creating an OSOPER Group (Optional)

Create an OSOPER group only if you want to identify a group of UNIX users with a limited set of database administrative privileges (SYSOPER operator privileges). For most installations, it is sufficient to create only the OSDBA group. If you want to use an OSOPER group, you must create it in the following circumstances:

- If an OSOPER group does not exist, for example, if this is the first installation of Oracle Database software on the system
- If an OSOPER group exists, but you want to give a different group of UNIX users database operator privileges in a new Oracle installation

If you require a new OSOPER group, create it by entering the following command, using the group name `oper` unless a group with that name already exists:

```
# /usr/sbin/groupadd oper
```

Creating the Oracle Software Owner User

You must create an Oracle software owner user in the following circumstances:

- If an Oracle software owner user does not exist, for example, if this is the first installation of Oracle software on the system
- If an Oracle software owner user exists, but you want to use a different UNIX user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation

Determine Whether an Existing Oracle Software Owner User Exists

To determine whether an Oracle software owner user named `oracle` exists, enter the following command:

```
# id oracle
```

If the `oracle` user exists, the output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

If the user exists, determine whether you want to use the existing user or create a new user. If you want to use the existing user, ensure that the user's primary group is the Oracle Inventory group and that it is a member of the appropriate OSDBA and OSOPER groups. See one of the following sections for more information:

Note: If necessary, contact your system administrator before using or modifying an existing user.

- If you want to use the existing Oracle software owner user, and the user's primary group is the Oracle Inventory group, see ["Verifying That the UNIX User nobody Exists"](#) on page 4-12
- To modify an existing user, see ["Modify an Existing Oracle Software Owner User"](#) on page 4-12.
- To create a new user, see the following section.

Create a New Oracle Software Owner User

If the Oracle software owner user does not exist or if you require a new Oracle software owner user, create it as follows, using the user name `oracle` unless a user with that name already exists.

1. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example, `oinstall`

- The `-G` option specifies the secondary groups, which must include the `OSDBA` group and if required, the `OSOPER` group, for example `dba` or `dba,oper`
2. Set the password of the `oracle` user:

```
# passwd oracle
```

To continue, refer to "[Verifying That the UNIX User nobody Exists](#)" on page 4-12.

Modify an Existing Oracle Software Owner User

If the `oracle` user exists, but its primary group is not `oinstall` or it is not a member of the appropriate `OSDBA` or `OSOPER` groups, you can modify it as follows:

- Enter a command similar to the following, specifying the primary group using the `-g` option and any required secondary groups using the `-G` option:

```
# /usr/sbin/usermod -g oinstall -G dba[,oper] oracle
```

Verifying That the UNIX User nobody Exists

Before installing the software, verify that the UNIX user `nobody` exists on the system:

1. To determine whether the user exists, enter the following command:

```
# id nobody
```

If this command displays information about the `nobody` user, you do not have to create that user.

2. If the `nobody` user does not exist, enter the following command to create it:

```
# /usr/sbin/useradd nobody
```

3. Repeat this procedure on all of the other cluster nodes.

Create Identical Users and Groups on Other Cluster Nodes

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

The Oracle software owner user and the Oracle Inventory, `OSDBA`, and `OSOPER` groups must exist and be identical on all cluster nodes. To create these identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, then create the user and groups with the same name and ID on the other cluster nodes.

Identifying the User and Group IDs

To determine the user ID (UID) of the Oracle software owner user and the group IDs (GID) of the Oracle Inventory, `OSDBA`, and `OSOPER` groups, follow these steps:

1. Enter the following command:

```
# id oracle
```

The output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

- From the output, identify the UID for the `oracle` user and the GIDs for the groups to which it belongs.

Create the User and Groups on the Other Cluster Nodes

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

- Log in to the next cluster node as `root`.
- Enter commands similar to the following to create the `oinstall` and `dba` groups, and if required, the `oper` group. Use the `-g` option to specify the correct GID for each group.

```
# /usr/sbin/groupadd -g 200 oinstall
# /usr/sbin/groupadd -g 201 dba
# /usr/sbin/groupadd -g 202 oper
```

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

- To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -u 200 -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-u` option specifies the user ID, which must be the user ID that you identified in the previous subsection.
- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`.
- The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba,oper`.

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

- Set the password of the `oracle` user:

```
# passwd oracle
```

Set Up User Equivalence for rsh and rcp on All Cluster Nodes

Note: This section describes how to set up user equivalence for `rcp`, which the Installer uses when copying Oracle software to the other cluster nodes. If you prefer, you can configure the Secure Shell (SSH) tool suite, so that the Installer uses `scp` instead of `rcp`. See the SSH documentation for information about setting up user equivalence for `scp`.

Before you install and use Oracle Real Application clusters, you must set up user equivalence on all cluster nodes, as follows:

1. On all cluster nodes, edit the `/etc/hosts.equiv` file in any text editor and add entries similar to the following for all nodes in the cluster:

```
node1-public oracle
node1-public.domain oracle
node2-public oracle
node2-public.domain oracle
.
.
.
```

In this example:

- `noden_public` is the host name associated with the primary public network interface
- `domain` is the domain name of the system, for example `mydomain.com`

Note: Alternatively, you can create a `.rhosts` file using the same format in the `oracle` user's home directory on all nodes in the cluster.

2. To test user equivalence on all cluster nodes, including the first node, enter commands similar to the following from the node where you will run the Installer:

```
# su oracle
$ rsh nodename1 date
$ rsh nodename2 date
.
.
.
$ exit
```

Each command should return only the date setting on the specified node. If a node prompts for a password or displays any other text, then user equivalence is not set up correctly on that node. In this case, verify that the `/etc/hosts.equiv` (or `~oracle/.rhosts`) file has the correct entries.

Configure Kernel Subsystem Attributes

Notes: The kernel subsystem attributes and shell limit values shown in the following section are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel subsystem attributes.

You must set the kernel subsystem attributes and shell limits on all cluster nodes.

Verify that the kernel subsystem attributes shown in the following table are set to values greater than or equal to the recommended value shown. The procedure following the table describes how to verify and set the values.

Subsystem	Attribute	Recommended Value
ipc	shm_max	4278190080 (4 GB minus 16 MB)
	shm_min	1
	shm_mni	256
	shm_seg	256
	ssm_threshold	Set this attribute to 0 only if the rad_gh_regions[n] or gh_chunks attributes are set in the vm subsystem. Otherwise, do not change the value.
proc	per_proc_stack_size	8388608 (8 MB) ¹
	max_per_proc_stack_size	33554432 (32 MB) ¹
	per_proc_data_size	335544320 (320MB)
	max_per_proc_data_size	335544320 (320MB)
	max_per_proc_address_space	Equal to the size of RAM or 1073741824 (1 GB), whichever is larger.
	per_proc_address_space	Equal to the size of RAM or 1073741824 (1 GB), whichever is larger.
vfs	fifo_do_adaptive	0
vm	new_wire_method	0
rdg	msg_size	32768
	max_objs	5120
	max_async_req	256
	max_sessions	500 (or at least 20 plus the value of the PROCESSES initialization parameter for all databases on the system, if this value is higher.)
	rdg_max_auto_msg_wires	0
	rdg_auto_msg_wires	0
	rt	aio_task_max_num

¹ The maximum value supported by Oracle for this parameter is 512 MB.

Note: If the current value for any kernel subsystem attribute is higher than the value listed in this table, *except* for the attributes whose recommended value is 0, do not change the value of that attribute.

To view the current value specified for these kernel subsystem attributes, and to change them if necessary, follow these steps:

1. To view the current values of the subsystem attributes, enter commands similar to the following:

```
# /sbin/sysconfig -q subsystem
```

For example, to view attribute values for the ipc subsystem, enter the following command:

```
# /sbin/sysconfig -q ipc
```

2. If you must change any of the current values, follow these steps:

- a. Create a backup copy of the `/etc/sysconfigtab` file, for example:

```
# cp /etc/sysconfigtab /etc/sysconfigtab.orig
```

- b. Using any text editor, create a file similar to the following, specifying the subsystems and attributes that you want to modify:

```
ipc:
    shm_max = 4278190080
    shm_min = 1
    shm_mni = 256
    shm_seg = 128

proc:
    per_proc_stack_size = 8388608
    max_per_proc_stack_size = 33554432
    per_proc_data_size = 201326592
    max_per_proc_data_size = 201326592
    max_per_proc_address_space = 4294967296
    per_proc_address_space = 4294967296
```

- c. Enter a command similar to the following to add the subsystem attributes to the `/etc/sysconfigtab` file:

```
# /sbin/sysconfigdb -m -f filename
```

In this example, *filename* is the name of the file you created in step b.

- d. Enter the following command to reboot the system:

```
# /sbin/shutdown -r now
```

- e. When the system restarts, log in and switch user to `root`.

3. Repeat this procedure on all other cluster nodes.

Identify Required Software Directories

You must identify or create four directories for the Oracle software, as follows:

- Oracle base directory

- Oracle Inventory directory
- CRS home directory
- Oracle home directory

The following subsections describe the requirements for these directories.

Oracle Base Directory

The Oracle base directory acts as a top-level directory for Oracle software installations. It is analogous to the `C:\Oracle` directory used for Oracle software installations on Windows systems. On UNIX systems, the Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to the following for the Oracle base directory:

```
/mount_point/app/oracle_sw_owner
```

In this example:

- *mount_point* is the mount point directory for the file system that will contain the Oracle software.

The examples in this guide use `/u01` for the mount point directory. However, you could choose another mount point directory, `/oracle` or `/opt/oracle` for example.

- *oracle_sw_owner* is the UNIX user name of the Oracle software owner, for example `oracle`.

You can use the same Oracle base directory for more than one installation or you can create separate Oracle base directories for different installations. If different UNIX users install Oracle software on the same system, each user must create a separate Oracle base directory. The following example Oracle base directories could all exist on the same system:

```
/u01/app/oracle
/u01/app/orauser
/opt/oracle/app/oracle
```

The following sections describe how to identify existing Oracle base directories that might be suitable for your installation and how to create a new Oracle base directory if necessary.

Regardless of whether you create a new Oracle base directory or decide to use an existing one, you must set the `ORACLE_BASE` environment variable to specify the full path to this directory.

Note: The Oracle base directory can be on a local file system, or on a supported cluster file system.

Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system. It is required by, and shared by, all Oracle software installations on a single system. The first time you install Oracle software on a system, the Installer prompts you to specify the path to this directory. If you are installing the software on a local file system, or on a Tru64 UNIX Cluster File System, then Oracle recommends that you choose the following path:

```
oracle_base/oraInventory
```

The Installer creates the directory that you specify and sets the correct owner, group, and permissions on it. You do not need to create it.

Note: All Oracle software installations rely on this directory. Make sure that you back it up regularly.

Do not delete this directory unless you have completely removed all Oracle software from the system.

CRS Home Directory

The CRS home directory is the directory where you choose to install the software for Oracle Cluster Ready Services. You must install CRS in a separate home directory. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
oracle_base/product/10.1.0/crs
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Oracle Home Directory

The Oracle home directory is the directory where you choose to install the software for a particular Oracle product. You must install different Oracle products, or different releases of the same Oracle product, in separate Oracle home directories. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the Oracle home directory:

```
oracle_base/product/10.1.0/db_1
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Identify or Create an Oracle Base Directory

Before starting the installation, you must either identify an existing Oracle base directory or if required, create a new one. This section contains information about the following:

- [Identifying an Existing Oracle Base Directory](#)
- [Creating a New Oracle Base Directory](#)

Note: You can choose to create a new Oracle base directory, even if other Oracle base directories exist on the system.

Identifying an Existing Oracle Base Directory

Existing Oracle base directories might not have paths that comply with OFA guidelines. However, if you identify an existing Oracle Inventory directory or existing Oracle home directories, you can usually identify the Oracle base directories, as follows:

- Identify an existing Oracle Inventory directory

Enter the following command on all cluster nodes to view the contents of the `oraInst.loc` file:

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

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inventory_loc` parameter identifies the Oracle Inventory directory (`oraInventory`) on that system. The parent directory of the `oraInventory` directory is typically an Oracle base directory. In the previous example, `/u01/app/oracle` is an Oracle base directory.

- Identify existing Oracle home directories

Enter the following command on all cluster nodes to view the contents of the `oratab` file:

```
# more /etc/oratab
```

If the `oratab` file exists, it contains lines similar to the following:

```
*:/u03/app/oracle/product/10.1.0/db_1:N
*/opt/orauser/infra_904:N
*/oracle/9.2.0:N
```

The directory paths specified on each line identify Oracle home directories. Directory paths that end with the user name of the Oracle software owner that you want to use are valid choices for an Oracle base directory. If you intend to use the `oracle` user to install the software, you could choose one of the following directories from the previous example:

```
/u03/app/oracle
/oracle
```

Note: If possible, choose a directory path similar to the first (`/u03/app/oracle`). This path complies with the OFA guidelines.

Before deciding to use an existing Oracle base directory for this installation, make sure that it satisfies the following conditions:

- It should not be on the same file system as the operating system.
- It must have an identical path on all cluster nodes or it must be on a supported cluster file system.

If you are not using a cluster file system, create identical Oracle base directories on the other nodes.

- It must have an identical path on all cluster nodes or it must be on a supported cluster file system.
If you are not using a cluster file system, create identical Oracle base directories on the other nodes.
- It must have sufficient free disk space on all cluster nodes, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB

Use the following command to determine the free disk space on the file system where the Oracle base directory is located:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see ["Creating a New Oracle Base Directory"](#) on page 4-20.
- If an Oracle base directory exists and you want to use it, see the ["Configure Disk Storage for Oracle CRS and Database Files"](#) section on page 4-21.

When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify the directory you chose.

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the following section.

Creating a New Oracle Base Directory

Before you create a new Oracle base directory, you must identify an appropriate file system with sufficient free disk space, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB

To identify an appropriate file system, follow these steps:

1. Use the `df -k` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has appropriate free space.

Note: The file system can be a local file system or a supported cluster file system.

The path to the Oracle base directory must be the same on all nodes.

3. Note the name of the mount point directory for the file system that you identified.

To create the Oracle base directory and specify the correct owner, group, and permissions for it, follow these steps:

1. Enter commands similar to the following to create the recommended subdirectories in the mount point directory that you identified and set the appropriate owner, group, and permissions on them:

```
# mkdir -p /mount_point/app/oracle_sw_owner
# chown -R oracle:oinstall /mount_point/app
# chmod -R 775 /mount_point/app
```

If the mount point you identified is `/u01` and `oracle` is the user name of the Oracle software owner, the recommended Oracle base directory path is as follows:

```
/u01/app/oracle
```

2. If necessary, repeat the commands listed in the previous step to create the same directory on the other cluster nodes.
3. When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify this directory.

Configure Disk Storage for Oracle CRS and Database Files

This section describes the storage options available to you for storing Oracle Cluster Ready Services files, Oracle Database Files, and optionally, Oracle Database recovery files. After you choose the storage method that you want to use for each file type, see the following sections for information about configuring the required storage.

Note: You do not have to use the same storage option for each type of file.

Choosing a Storage Option for Oracle CRS Files

Before you install Oracle Cluster Ready Services, you must choose the storage option that you want to use for the Oracle Cluster Registry (100 MB) and CRS voting disk (20 MB). You cannot use Automatic Storage Management to store these files, because they must be accessible before any Oracle instance starts.

Apart from ASM, you can use any of the storage options listed in the next section for the Oracle CRS files.

Choosing a Storage Option for Oracle Database Files

If you want to create a database during the installation, you must choose one of the following storage options for the database files:

- Cluster File System with TruCluster V5.1B
- Automatic Storage Management
- Logical Storage Manager with TruCluster V5.1B
- Raw partitions

Note: For the most up-to-date information about supported storage options for RAC installations, see the Certify pages on the Oracle *MetaLink* Web site:

<http://metalink.oracle.com>

Choosing a Storage Option for Oracle Database Recovery Files

Additionally, if you want to enable automated backups during the installation, you must choose one of the following storage options for recovery files (the flash recovery area):

- Cluster File System with TruCluster V5.1B
- Automatic Storage Management

The storage option that you choose for recovery files can be the same as or different to the option you choose for the database files.

Configuring Disk Storage

For information about how to configure disk storage before you start the installation, see one of the following sections depending on your choice:

- To use a file system for Oracle CRS, database, or recovery file storage, see "[Create Directories for Oracle CRS, Database, or Recovery Files](#)" on page 4-22.
- To use ASM for database or recovery file storage, see "[Configure Disks for Automatic Storage Management](#)" on page 4-25.

Create Directories for Oracle CRS, Database, or Recovery Files

If you decide to place the Oracle CRS, database, or recovery files on a file system, use the following guidelines when deciding where to place them.

Guidelines for Placing Oracle CRS Files on a File System

The Installer does not suggest a default location for the Oracle Cluster Registry (OCR) or the Oracle CRS voting disk. If you choose to create these files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk.

Note: NAS storage is currently supported only if you are using Fujitsu PRIMECLUSTER and a certified NAS device.

- It must have at least 100 MB of free disk space for the OCR and 20 MB of free disk space for the CRS voting disk.
- For improved reliability, you should choose a file system on a highly available storage device, for example, a RAID device that implements mirroring.
- If you are placing the Oracle Cluster Ready Services software on a shared file system, you can use the same file system for these files.
- The `oracle` user must have write permissions to create the files in the path you specify.

Note: If you are upgrading from Oracle9i Release 2, you can continue to use the raw device or shared file that you used for the SRVM configuration repository instead of creating a new file for the OCR.

Guidelines for Placing Oracle Database Files on a File System

If you choose to place the Oracle database files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk.

Note: NAS storage is currently supported only if you are using Fujitsu PRIMECLUSTER and a certified NAS device.

- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.
This default location is not recommended for production databases.
- You can choose either a single file system or more than one file system to store the database files:
 - If you want to use a single file system, choose a file system on a physical device that is dedicated to the database.
For best performance and reliability, choose a RAID device or a logical volume on more than one physical device and implement the stripe-and-mirror-everything (SAME) methodology.
 - If you want to use more than one file system, choose file systems on separate physical devices that are dedicated to the database.
This method enables you to distribute physical I/O and create separate control files on different devices for increased reliability. It also enables you to fully implement the OFA guidelines. You must choose either the Advanced database creation option or the Custom installation type during the installation to implement this method.
- If you intend to create a preconfigured database during the installation, the file system (or file systems) that you choose must have at least 1.2 GB of free disk space.
For production databases, you must estimate the disk space requirement depending on the use you want to make of the database.
- For optimum performance, the file systems that you choose should be on physical devices that are used only by the database.
- The `oracle` user must have write permissions to create the files in the path you specify.

Guidelines for Placing Oracle Recovery Files on a File System

Note: You must choose a location for recovery files only if you intend to enable automated backups during the installation.

If you choose to place the Oracle recovery files on a file system, use the following guidelines when deciding where to place them:

- To prevent disk failure from making both the database files and the recovery files unavailable, place the recovery files in a file system on a different physical disk from the database files.

Note: Alternatively use an ASM disk group with a normal or high redundancy level for either or both file types.

- You must choose a shared file system, for example, a cluster file system on a shared disk.

Note: NAS storage is currently supported only if you are using Fujitsu PRIMECLUSTER and a certified NAS device.

- The file system that you choose should have at least 2 GB of free disk space.

The disk space requirement is the default disk quota configured for the flash recovery area (specified by the `DB_RECOVERY_FILE_DEST_SIZE` initialization parameter).

If you choose the Custom installation type or the Advanced database configuration option, you can specify a different disk quota value. After you create the database, you can also use Oracle Enterprise Manager Grid Control or Database Control to specify a different value.

For more information about sizing the flash recovery area, see the *Oracle Backup and Recovery Basics* manual.

- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.

This default location is not recommended for production databases.

- The `oracle` user must have write permissions to create the files in the path you specify.

Creating Required Directories

Note: You must complete this procedure only if you want to place the Oracle CRS, database, or recovery files on a separate file system to the Oracle base directory.

To create directories for the Oracle CRS, database, or recovery files on separate file systems to the Oracle base directory, follow these steps:

1. If necessary, configure the shared file systems that you want to use and mount them on each node.

Note: Make sure that the file systems are configured to mount automatically when a node reboots.

2. Use the `df -k` command to determine the free disk space on each mounted file system.
3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
CRS files	Choose a file system with at least 120 MB of free disk space
Datafiles	Choose either: <ul style="list-style-type: none"> ■ A single file system with at least 1.2 GB of free disk space ■ Two or more file systems with at least 1.2 GB of free disk space in total
Recovery files	Choose a file system with at least 2 GB of free disk space.

If you are using the same file system for more than one type of file, add the disk space requirements for each type to determine the total disk space requirement.

4. Note the names of the mount point directories for the file systems that you identified.
5. Enter commands similar to the following to create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:
 - CRS file directory:

```
# mkdir /mount_point/oracrs
# chown oracle:oinstall /mount_point/oracrs
# chmod 775 /mount_point/oracrs
```
 - Database file directory:

```
# mkdir /mount_point/oradata
# chown oracle:oinstall /mount_point/oradata
# chmod 775 /mount_point/oradata
```
 - Recovery file directory (flash recovery area):

```
# mkdir /mount_point/flash_recovery_area
# chown oracle:oinstall /mount_point/flash_recovery_area
# chmod 775 /mount_point/flash_recovery_area
```
6. To use ASM for storage, then refer to the following section, "[Configure Disks for Automatic Storage Management](#)".

Otherwise see "[Verify that the Cluster Software is Running](#)" on page 4-31.

Configure Disks for Automatic Storage Management

This section describes how to configure disks for use with ASM. Before you configure the disks, you must determine the number of disks and the amount of free disk space

that you require. The following sections describe how to identify the requirements and configure the disks:

- [Identifying Storage Requirements for ASM](#)
- [Using an Existing ASM Disk Group](#)
- [Using an Existing ASM Disk Group](#)

Note: Although this section refers to disks, you can also use zero-padded files on a certified NAS storage device in an ASM disk group. See the appendices in *Oracle Database Installation Guide for UNIX Systems* for information about creating and configuring NAS-based files for use in an ASM disk group.

Identifying Storage Requirements for ASM

To identify the storage requirements for using ASM, you must determine how many devices and the amount of free disk space that you require. To complete this task, follow these steps:

1. Determine whether you want to use ASM for Oracle database files (datafiles), recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. One can use the file system, while the other uses ASM.

If you choose to enable automated backups and you do not have a cluster file system available, you must choose ASM for recovery file storage.

If you enable automated backups during the installation, you can choose ASM as the storage mechanism for recovery files by specifying an ASM disk group for the flash recovery area. Depending how you choose to create a database during the installation, you have the following options:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to use the same ASM disk group for datafiles and recovery files, or you can choose to use different disk groups for each file type.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must use the same ASM disk group for database files and recovery files.
2. Choose the ASM redundancy level that you want to use for the ASM disk group.

The redundancy level that you choose for the ASM disk group determines how ASM mirrors files in the disk group and determines the number of disks and amount of disk space that you require, as follows:

- 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 ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use only RAID or similar devices that provide their own data protection mechanisms as disk devices in this type of disk group.

- Normal redundancy

In a normal redundancy disk group, ASM uses two-way mirroring by default, to increase performance and reliability. 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 most installations, Oracle recommends that you use normal redundancy disk groups.

- High redundancy

In a high redundancy disk group, 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.

While high redundancy disk groups do provide a high level of data protection, you must consider the higher cost of additional storage devices before deciding to use this redundancy level.

3. Determine the total amount of disk space that you require for the datafiles and recovery files.

Use the following table to determine the minimum number of disks and the minimum disk space requirements for the installation:

Redundancy Level	Minimum Number of Disks	Datafiles	Recovery Files	Both File Types
External	1	1.15 GB	2.3 GB	3.45 GB
Normal	2	2.3 GB	4.6 GB	6.9 GB
High	3	3.45 GB	6.9 GB	10.35 GB

You must also add additional disk space for the ASM metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

$$15 + (2 * \text{number_of_disks}) + (126 * \text{number_of_ASM_instances})$$

For example, for a four-node RAC installation, using three disks in a high redundancy disk group, you require an additional 525 MB of disk space:

$$(15 + (2 * 3) + (126 * 4)) = 525$$

If an existing ASM instance exists on the system, 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 installation.

The following section describes how to identify existing disk groups and determine the free disk space that they contain.

4. Optionally identify failure groups for the ASM disk group devices.

Note: You need to complete this step only if you intend to use an installation method that runs DBCA in interactive mode, for example, if you intend to choose the Custom installation type or the Advanced database configuration option. Other installation types do not enable you to specify failure groups.

If you intend to use a normal or high redundancy disk group, 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 SCSI controller, the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To avoid failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

5. If you are sure that a suitable disk group does not exist on the system, 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 ASM disk group should be the same size and have the same performance characteristics.
 - Do not specify more than one partition on a single physical disk as a disk group device. 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 ASM disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing ASM from optimizing I/O across the physical devices.

For information about completing this task, see "[Configure Disks for Automatic Storage Management](#)" on page 4-25.

Using an Existing ASM Disk Group

If you want to use ASM as the storage option for either database or recovery files, and an existing ASM disk group exists, you have the following choices, depending on the installation method that you select:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to create a new disk group or use an existing one.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must choose an existing disk group for the new database; you cannot create a new

disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The ASM instance that manages the existing disk group can be running in a different Oracle home directory.

To determine whether an existing ASM disk group exists, or to determine whether there is sufficient disk space in a disk group, you can use Oracle Enterprise Manager Database Grid Control or Database Control. Alternatively, you can use the following procedure:

1. View the contents of the `oratab` file to determine whether an ASM instance is configured on the system:

```
# more /etc/oratab
```

If an ASM instance is configured on the system, the `oratab` file should contain a line similar to the following:

```
+ASM:oracle_home_path:N
```

In this example, `+ASM` is the system identifier (SID) of the ASM instance and `oracle_home_path` is the Oracle home directory where it is installed. By convention, the SID for an ASM instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the ASM instance that you want to use.
3. Connect to the ASM instance as the `SYS` user with `SYSDBA` privilege and start the instance if necessary:

```
# $ORACLE_HOME/bin/sqlplus "SYS/SYS_password as SYSDBA"
SQL> STARTUP
```

4. Enter the following command to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

```
SQL> SELECT NAME,TYPE,TOTAL_MB,FREE_MB FROM V$ASM_DISKGROUP;
```

5. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.
6. If necessary, install or identify the additional disk devices required to meet the storage requirements listed in the previous section.

Note: If you are adding devices to an existing disk group, Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Configuring Disks for ASM

To configure disks for use with ASM on Tru64 UNIX, follow these steps:

1. If necessary, install the shared disks that you intend to use for the disk group and reboot the system.
2. To make sure that the disks are available, enter the following command:

```
# file /dev/rdisk/dsk*c
```

The output from this command is similar to the following:

```
/dev/rdisk/dsk0c:      character special (19/38) SCSI #1 "RZ1CB-CA" disk #2
                      (SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk10c:    character special (19/218) SCSI #2 "RZ1DB-CS" disk #4
                      (SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk11c:    character special (19/234) SCSI #2 "RZ1DB-CS" disk #5
                      (SCSI ID #3) (SCSI LUN #0)
/dev/rdisk/dsk12c:    character special (19/250) SCSI #2 "RZ1DB-CS" disk #6
                      (SCSI ID #4) (SCSI LUN #0)
```

This command displays information about each disk attached to the system, including the raw character device name for partition *c*, which represents the entire disk (*/dev/rdisk/dsknc*).

- From the list, identify the device names for the disk devices that you want to add to a disk group.
- To determine whether the disk is unused, enter a command similar to the following:

```
# /sbin/disklabel dskn
```

This command displays information about the partitions on the disk, and identifies whether they are in use.

- Enter the following command to verify that the device you identified is not used in an AdvFS domain:

```
# ls /etc/fdmns/*/*
```

This command displays information about the devices included in AdvFS domains in the following format, identifying the domains and the disks associated with each domain:

```
/etc/fdmns/domain_name/dsknp
```

In this example, *domain_name* is the name of the domain, *n* is the disk number, and *p* is the partition letter.

- Enter the following command to verify that the device is not mounted as a UFS file system:

```
# /sbin/mount
```

The device you want to add to the ASM disk group must not be listed in the output.

- Enter the following command to verify that the device you identified is not already part of a Tru64 UNIX Logical Storage Manager (LSM) disk group:

```
# /sbin/voldisk list
```

Disks listed without DISK and GROUP names, and with a status of unknown are not LSM disks.

- Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for each disk that you want to add to a disk group:

```
# chown oracle:dba /dev/rdisk/rdsknc
# chmod 660 /dev/rdisk/rdsknc
```

In this example, the device name specifies partition `c`, which represents the entire disk.

Note: If you are using a multi-pathing disk driver with ASM, ensure that you set the permissions only on the correct logical device name for the disk.

Verify that the Cluster Software is Running

Note: The cluster software mentioned in the following sections is not required to install Oracle Real Application Clusters. However, if this software is installed, Oracle Cluster Ready Services (CRS) can integrate with it.

To verify that the TruCluster software is running, enter the following command:

```
# /usr/sbin/clu_get_info
```

Make sure that the correct number of nodes are configured and that all cluster nodes are up.

If necessary, see the TruCluster documentation for information about troubleshooting the cluster.

Stop Existing Oracle Processes

Caution: If you are installing additional Oracle Database 10g products in an existing Oracle home, stop *all* processes running in the Oracle home. You must complete this task to enable the Installer to relink certain executables and libraries.

If you choose to create a database during the installation, most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the IPC key value EXTPROC. However, if an existing Oracle Net listener process is using the same port or key value, the Installer can only configure the new listener; it cannot start it. To ensure that the new listener process starts during the installation, you must shut down any existing listeners before starting the Installer.

To determine whether an existing listener process is running and to shut it down if necessary, follow these steps:

1. Switch user to `oracle`:

```
# su - oracle
```

2. Enter the following command to determine whether a listener process is running and to identify its name and the Oracle home directory in which it is installed:

```
$ ps -ef | grep tnslnsr
```

This command displays information about the Oracle Net listeners running on the system:

```
... oracle_home1/bin/tnslsnr LISTENER -inherit
```

In this example, *oracle_home1* is the Oracle home directory where the listener is installed and LISTENER is the listener name.

Note: If no Oracle Net listeners are running, see the "[Configure the oracle User's Environment](#)" section on page 4-32 to continue.

3. Set the ORACLE_HOME environment variable to specify the appropriate Oracle home directory for the listener:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=oracle_home1
$ export ORACLE_HOME
```

- C or tcsh shell:

```
$ setenv ORACLE_HOME oracle_home1
```

4. Enter the following command to identify the TCP/IP port number and IPC key value that the listener is using:

```
$ $ORACLE_HOME/bin/lsnrctl status listenername
```

Note: If the listener uses the default name LISTENER, you do not have to specify the listener name in this command.

5. Enter a command similar to the following to stop the listener process:

```
$ $ORACLE_HOME/bin/lsnrctl stop listenername
```

6. Repeat this procedure to stop all listeners running on this system and on all other cluster nodes.

Configure the oracle User's Environment

You run the Installer from the *oracle* account. However, before you start the Installer you must configure the environment of the *oracle* user. To configure the environment, you must:

- Set the default file mode creation mask (*umask*) to 022 in the shell startup file.
- Set the DISPLAY and ORACLE_BASE environment variables.

To set the *oracle* user's environment, follow these steps:

1. Start a new terminal session, for example, an X terminal (*xterm*).
2. Enter the following command to ensure that X Window applications can display on this system:

```
$ xhost +
```

3. If you are not already logged in to the system where you want to install the software, log in to that system as the *oracle* user.
4. If you are not logged in as the *oracle* user, switch user to *oracle*:

```
$ su - oracle
```

5. To determine the default shell for the `oracle` user, enter the following command:

```
$ echo $SHELL
```

6. Open the `oracle` user's shell startup file in any text editor:

- Bourne shell (`sh`), Bash shell (`bash`), or Korn shell (`ksh`):

```
$ vi .profile
```

- C shell (`cs`h or `tcsh`):

```
% vi .login
```

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

```
umask 022
```

8. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, remove the appropriate lines from the file.

9. Save the file and exit from the editor.

10. To run the shell startup script, enter one of the following commands:

- Bourne, Bash, or Korn shell:

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

- C shell:

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

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

- Bourne, Bash, or Korn shell:

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

- C shell:

```
$ setenv DISPLAY local_host:0.0
```

In this example, `local_host` is the host name or IP address of the system you want to use to display the Installer (your workstation or PC).

12. If you determined that the `/tmp` directory has less than 400 MB of free disk space, identify a file system with at least 400 MB of free space and set the `TEMP` and `TMPDIR` environment variables to specify a temporary directory on this file system:

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

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

```
$ su - root
# mkdir /mount_point/tmp
# chmod a+wr /mount_point/tmp
# exit
```

- c. Enter commands similar to the following to set the TEMP and TMPDIR environment variables:

- * Bourne, Bash, or Korn shell:

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

- * C shell:

```
$ setenv TEMP /mount_point/tmp
$ setenv TMPDIR /mount_point/tmp
```

13. Enter commands similar to the following to set the ORACLE_BASE environment variable:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle
$ export ORACLE_BASE
```

- C shell:

```
% setenv ORACLE_BASE /u01/app/oracle
```

In these examples, /u01/app/oracle is the Oracle base directory that you created earlier.

14. If you are using raw devices for database storage, set the DBCA_RAW_CONFIG environment variable to specify the full path to the raw device mapping file:

- Bourne, Bash, or Korn shell:

```
$ DBCA_RAW_CONFIG=$ORACLE_BASE/oradata/dbname/dbname_raw.conf
$ export DBCA_RAW_CONFIG
```

- C shell:

```
% setenv DBCA_RAW_CONFIG=$ORACLE_BASE/oradata/dbname/dbname_raw.conf
```

15. Enter the following command to ensure that the ORACLE_HOME and TNS_ADMIN environment variables are not set:

```
$ unset ORACLE_HOME
$ unset TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value it specifies as the default path for the Oracle home directory. However, if you set the ORACLE_BASE environment variable, Oracle recommends that you unset the ORACLE_HOME environment variable and choose the default path suggested by the Installer.

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

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

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

Pre-Installation Tasks for Installing RAC on Linux-Based Systems

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer. It includes information about the following tasks:

- Log In to the System as root
- Check the Hardware Requirements
- Check the Network Requirements
- Check the Software Requirements
- Create Required UNIX Groups and User
- Configure Kernel Parameters and Shell Limits
- Identify Required Software Directories
- Identify or Create an Oracle Base Directory
- Configure Disk Storage for Oracle CRS and Database Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Configure Raw Partitions
- Verify that the Required Software is Running
- Stop Existing Oracle Processes
- Configure the oracle User's Environment

Log In to the System as root

Before you install the Oracle software, you must complete several tasks as the `root` user. To log in as the `root` user, complete one of the following procedures:

Note: Unless you intend to complete a silent installation, you must install the software from an X Window System workstation, an X terminal, or a PC or other system with X server software installed.

For more information about non-interactive and silent installations, see the appendices in *Oracle Database Installation Guide for UNIX Systems*.

- If you are installing the software from an X Window System workstation or X terminal:
 1. Start a local terminal session, for example, an X terminal (`xterm`).
 2. If you are not installing the software on the local system, enter the following command to enable remote hosts to display X applications on the local X server:

```
$ xhost +
```
 3. If you want to install the software on a remote system, enter a command similar to the following to connect to that system:

```
$ telnet remote_host
```
 4. If you are not logged in as the `root` user, enter the following command to switch user to `root`:

```
$ su - root
password:
#
```
- If you are installing the software from a PC or other system with X server software installed:

Note: If necessary, see your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

1. Start the X server software.
2. Configure the security settings of the X server software to permit remote hosts to display X applications on the local system.
3. Connect to the remote system where you want to install the software and start a terminal session on that system, for example, an X terminal (`xterm`).
4. If you are not logged in as the `root` user on the remote system, enter the following command to switch user to `root`:

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

Check the Hardware Requirements

The system must meet the following minimum hardware requirements:

- 512 MB of physical RAM
- 1 GB of swap space (or twice the size of RAM)
On systems with 2 GB or more of RAM, the swap space can be between one and two times the size of RAM.
- 400 MB of disk space in the `/tmp` directory
- Up to 4 GB of disk space for the Oracle software, depending on the installation type
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirements for databases that use Automatic Storage Management (ASM) or raw device storage is described later in this chapter.

Additional disk space, either on a file system or in an ASM disk group, is required for the flash recovery area if you choose to configure automated backups.

To ensure that the system meets these requirements, follow these steps:

1. To determine the physical RAM size, enter the following command:

```
# grep MemTotal /proc/meminfo
```

If the size of the physical RAM installed in the system is less than the required size, you must install more memory before continuing.
2. To determine the size of the configured swap space, enter the following command:

```
# grep SwapTotal /proc/meminfo
```

If necessary, see your operating system documentation for information about how to configure additional swap space.
3. To determine the amount of disk space available in the `/tmp` directory, enter the following command:

```
# df -k /tmp
```

If there is less than 400 MB of disk space available in the `/tmp` directory, complete one of the following steps:

 - Delete unnecessary files from the `/tmp` directory to achieve the required disk space.
 - Set the `TEMP` and `TMPDIR` environment variables when setting the `oracle` user's environment (described later).
 - Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.
4. To determine the amount of free disk space on the system, enter the following command:

```
# df -k
```

5. To determine whether the system architecture can run the software, enter the command:

```
# grep "model name" /proc/cpuinfo
```

This command displays the processor type. Verify that the processor architecture matches the Oracle software release that you want to install.

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

Check the Network Requirements

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Real Application Clusters installation.

Note: For the most up-to-date information about supported network protocols and hardware for RAC installations, see the *Certify* pages on the Oracle *MetaLink* Web site:

<http://metalink.oracle.com>

Network Hardware Requirements

Each node in the cluster must meet the following requirements:

- Each node must have at least two network adapters; one for the public network interface and one for the private network interface (the interconnect).
- The interface names associated with the network adapters for each network must be the same on all nodes.
- For increased reliability, you can configure redundant public and private network adapters for each node.
- For the public network, each network adapter must support TCP/IP.
- For the private network, the interconnect must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended).

Note: UDP is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS.

IP Address Requirements

Before starting the installation, you must identify or obtain the following IP addresses for each node:

- An IP address and an associated host name registered in the domain name service (DNS) for each public network interface
- One unused virtual IP address and an associated virtual host name registered in DNS that you will configure for the primary public network interface

The virtual IP address must be in the same subnet as the associated public interface. After installation, you can configure clients to use the virtual host name or IP address. If a node fails, its virtual IP address fails over to another node.

- A private IP address and optional host name for each private interface
- Oracle recommends that you use private network IP addresses for these interfaces, for example: 10.*.* or 192.168.*.*. You can use the `/etc/hosts` file on each node to associate private host names with private IP addresses.

For example, if each node has two public and two private interfaces, you might have the following host names and IP addresses on one of the nodes (`rac1`) and similar host names and IP addresses on the other nodes:

Host Name	Type	IP Address	Registered In
<code>rac1.mydomain.com</code>	Public	143.47.43.100	DNS
<code>rac1-2.mydomain.com</code>	Public	143.46.51.101	DNS
<code>rac1-vip.mydomain.com</code>	Virtual	143.46.43.104	DNS
<code>rac1-priv1</code>	Private	10.0.0.1	<code>/etc/hosts</code>
<code>rac1-priv2</code>	Private	10.0.0.2	<code>/etc/hosts</code>

Checking the Network Requirements

To verify that each node meets the requirements, follow these steps:

1. If necessary, install the network adapters for the public and private networks and configure them with either public or private IP addresses.
2. Register the host names and IP addresses for the public network interfaces in DNS.
3. For each node, register one virtual host name and IP address in DNS.
4. For each private interface on every node, add a line similar to the following to the `/etc/hosts` file on all nodes, specifying the private IP address and associated private host name:

```
10.0.0.1    rac1-priv1
```

5. To identify the interface name and associated IP address for every network adapter, enter the following command:

```
# /sbin/ifconfig
```

From the output, identify the interface name and IP address for all network adapters that you want to specify as public or private network interfaces.

Note: When you install Oracle CRS and RAC, you will require this information.

Check the Software Requirements

See the following sections for information about checking the software requirements.

- [Check for Required Software](#)
- [Check for Required Patches](#)

Note: The Oracle Universal Installer performs checks on your system to verify that it meets the requirements listed Linux. To ensure that these checks pass, verify the requirements before you start the Installer.

Check for Required Software

Depending on the products that you intend to install, verify that the following software is installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	<p>One of the following operating system versions:</p> <ul style="list-style-type: none"> ■ Red Hat Enterprise Linux AS/ES 2.1 or 3 (x86) ■ UnitedLinux 1.0, service pack 3 or higher (x86) <hr/> <p>The following packages (or higher versions) must be installed:</p> <p>All distributions:</p> <p>make-3.79 binutils-2.11.90.0.8-12</p> <p>Red Hat Enterprise Linux 2.1 (x86):</p> <p>gcc-2.96.108.1 openmotif-2.1.30-11</p> <p>Red Hat Enterprise Linux 3:</p> <p>gcc-3.2.3-2 compat-db-4.0.14.5 compat-gcc-7.3-2.96.122 compat-gcc-c++-7.3-2.96.122 compat-libstdc++-7.3-2.96.122 compat-libstdc++-devel-7.3-2.96.122 openmotif21-2.1.30-8 setarch-1.3-1</p> <p>UnitedLinux 1.0 (x86):</p> <p>gcc-3.2.2-38 openmotif-2.2.2-124</p>
Oracle Spatial	<p>X Window System development package:</p> <ul style="list-style-type: none"> ■ Red Hat: <ul style="list-style-type: none"> XFree86 XFree86-devel ■ UnitedLinux: <ul style="list-style-type: none"> xf86 xdevel <p>Note: This software is required only to build the sample programs.</p>
Oracle Messaging Gateways	<p>IBM WebSphere MQ V5.3, client and server:</p> <p>MQSeriesClient MQSeriesServer MQSeriesRuntime</p>

Installation Type or Product	Requirement
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	The version of Gnu gcc listed previously for your distribution is supported for use with these products. Note: Intel C++ Compiler v7.1.0.28 or higher is also supported, however it is not required for installation.
Oracle JDBC/OCI Drivers	You can use the following optional JDK version with the Oracle JDBC/OCI drivers, however it is not required for the installation: <ul style="list-style-type: none"> ■ Sun JDK 1.3.1_09 with the JNDI extension Note: JDK 1.4.2 is installed with this release.
Real Application Clusters	Oracle Cluster File System (OCFS) version 1.0.9-12 or later: Note: OCFS is required only if you want to use a cluster file system for database files. If you want to use Automatic Storage Management or raw devices for database file storage, you do not need to install OCFS. ocfs-support ocfs-tools ocfs-kernel_version

To ensure that the system meets these requirements, follow these steps:

1. To determine which distribution and version of Linux is installed, enter the following command:

```
# cat /etc/issue
```

Note: Only the distributions and versions listed in the previous table are supported. Do not install the software on other versions of Linux.

2. On UnitedLinux 1.0 systems only, enter the following command to determine whether Service Pack 3 is installed:

```
# uname -r
2.4.21-138-default
```

This command shows the kernel version (2.4.21-138) and type (default). If the kernel version is less than 2.4.21-138, contact your UnitedLinux vendor for information about obtaining and installing Service Pack 3.

3. To determine whether the required packages are installed, enter a command similar to the following:

```
# rpm -q package_name
```

If a package is not installed, install it from your Linux distribution media.

4. To determine whether OCFS is installed, enter the following command:

```
# rpm -qa | grep ocfs
```

If you want to install the Oracle database files on an OCFS file system and the packages are not installed, download them from the following Web site. Follow the instructions listed with the kit to install the packages and configure the file system:

<http://oss.oracle.com/projects/ocfs/>

Check for Required Patches

Depending on the products that you intend to install, verify that the following patches are installed on the system. The procedure following the table describes how to check these requirements.

Note: If the following patches are superseded by later versions, install the latest version.

Installation Type or Product	Requirement
All installations	Red Hat Enterprise Linux 2.1 (x86) errata: <ul style="list-style-type: none"> ▪ RHSA-2003:195-06 Updated kernel addresses security vulnerabilities ▪ RHSA-2003:022-09 Updated glibc packages fix vulnerabilities in resolver
Oracle Messaging Gateway	Corrective service diskette (CSD) for WebSphere MQ: CSD05 or later for WebSphere MQ V5.3

To ensure that the system meets these requirements, follow these steps:

1. On Red Hat Enterprise Linux 2.1 systems only, follow these steps:
 - a. To determine whether kernel update RHSA-2003:195-06 or later is installed, enter the following command:

```
# uname -r
2.4.9-e.25
```

This command shows the kernel version (2.4.9) and errata level (e.25). If the errata level is less than 25, install the latest kernel update. See the following Web site for information about kernel update RHSA-2003:195-06:

<https://rhn.redhat.com/errata/RHSA-2003-195.html>

- b. To determine whether fix RHSA-2003:022-09 is installed, enter the following command:

```
# rpm -q glibc
glibc-2.2.4-31
```

If the version shown is less than 2.2.4-31, see the following Web site for information about fix RHSA-2003:022-09:

<https://rhn.redhat.com/errata/RHSA-2003-022.html>

2. If you require a CSD for WebSphere MQ, see the following Web site for download and installation information:

<http://www.ibm.com/software/integration/mqfamily/support/summary/lin.html>

Create Required UNIX Groups and User

Depending on whether this is the first time Oracle software is being installed on this system and on the products that you are installing, you may need to create several UNIX groups and a UNIX user.

The following UNIX groups and user are required if you are installing Oracle Database:

- The OSDBA group (`dba`)

You must create this group the first time you install Oracle Database software on the system. It identifies UNIX users that have database administrative privileges (the SYSDBA privilege). The default name for this group is `dba`.

If you want to specify a group name other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of this group. In this case, the Installer prompts you to specify the name of this group.

- The OSOPER group (`oper`)

This is an optional group. Create this group if you want a separate group of UNIX users to have a limited set of database administrative privileges (the SYSOPER privilege). By default, members of the OSDBA group also have the SYSOPER privilege.

If you want to specify a separate OSOPER group, other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of the `dba` group. In this case, the Installer prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

- An unprivileged user (`nobody`)

You must verify that the unprivileged user `nobody` exists on the system. The `nobody` user must own the external jobs (`ext job`) executable after the installation.

The following UNIX group and user are required for all installations:

- The Oracle Inventory group (`oinstall`)

You must create this group the first time you install Oracle software on the system. The usual name chosen for this group is `oinstall`. This group owns the Oracle inventory, which is a catalog of all Oracle software installed on the system.

Note: If Oracle software is already installed on the system, the existing Oracle Inventory group must be the primary group of the UNIX user that you use to install new Oracle software. The following sections describe how to identify an existing Oracle Inventory group.

- The Oracle software owner user (`oracle`)

You must create this user the first time you install Oracle software on the system. This user owns all of the software installed during the installation. The usual name chosen for this user is `oracle`. This user must have the Oracle Inventory group as its primary group. It must also have the OSDBA and OSOPER groups as a secondary groups.

A single Oracle Inventory group is required for all installations of Oracle software on the system. After the first installation of Oracle software, you must use the same Oracle Inventory group for all subsequent Oracle software installations on that system. However, you can choose to create different Oracle software owner users, OSDBA groups, and OSOPER groups (other than `oracle`, `dba`, and `oper`) for separate installations. By using different groups for different installations, members of these different groups have DBA privileges only on the associated databases rather than on all databases on the system.

See Also: The *Oracle Database Administrator's Reference for UNIX Systems* and the *Oracle Database Administrator's Guide* contain more information about the OSDBA and OSOPER groups and the SYSDBA and SYSOPER privileges.

The following sections describe how to create the required UNIX user and groups.

Note: The following sections describe how to create local users and groups. As an alternative to creating local users and groups, you could create the appropriate users and groups in a directory service, for example, Network Information Services (NIS). For information about using directory services, contact your system administrator or see your operating system documentation.

Creating the Oracle Inventory Group

You must create the Oracle Inventory group if it does not already exist. The following subsections describe how to determine the Oracle Inventory group name, if it exists, and how to create it if necessary.

Determine Whether the Oracle Inventory Group Exists

When you install Oracle software on the system for the first time, the Installer creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group and the path of the Oracle Inventory directory. To determine whether the Oracle Inventory group exists, enter the following command:

```
# more /etc/oraInst.loc
```

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inst_group` parameter shows the name of the Oracle Inventory group (`oinstall`).

Create the Oracle Inventory Group

If the `oraInst.loc` file does not exist, create the Oracle Inventory group, `oinstall`, by entering the following command:

```
# /usr/sbin/groupadd oinstall
```

Creating the OSDBA Group

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system
- An OSBDA group exists, but you want to give a different group of UNIX users database administrative privileges in a new Oracle installation

If the OSDBA group does not exist or if you require a new OSDBA group, create it by entering the following command, using the group name `dba` unless a group with that name already exists:

```
# /usr/sbin/groupadd dba
```

Creating an OSOPER Group (Optional)

Create an OSOPER group only if you want to identify a group of UNIX users with a limited set of database administrative privileges (SYSOPER operator privileges). For most installations, it is sufficient to create only the OSDBA group. If you want to use an OSOPER group, you must create it in the following circumstances:

- If an OSOPER group does not exist, for example, if this is the first installation of Oracle Database software on the system
- If an OSOPER group exists, but you want to give a different group of UNIX users database operator privileges in a new Oracle installation

If you require an new OSOPER group, create it by entering the following command, using the group name `oper` unless a group with that name already exists:

```
# /usr/sbin/groupadd oper
```

Creating the Oracle Software Owner User

You must create an Oracle software owner user in the following circumstances:

- If an Oracle software owner user does not exist, for example, if this is the first installation of Oracle software on the system
- If an Oracle software owner user exists, but you want to use a different UNIX user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation

Determine Whether an Existing Oracle Software Owner User Exists

To determine whether an Oracle software owner user named `oracle` exists, enter the following command:

```
# id oracle
```

If the `oracle` user exists, the output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

If the user exists, determine whether you want to use the existing user or create a new user. If you want to use the existing user, ensure that the user's primary group is the Oracle Inventory group and that it is a member of the appropriate OSDBA and OSOPER groups. See one of the following sections for more information:

Note: If necessary, contact your system administrator before using or modifying an existing user.

- If you want to use the existing Oracle software owner user, and the user's primary group is the Oracle Inventory group, see the "[Verifying That the UNIX User nobody Exists](#)" section on page 5-12.
- To modify an existing user, see the "[Modify an Existing Oracle Software Owner User](#)" section on page 5-12.
- To create a new user, see the following section.

Create a New Oracle Software Owner User

If the Oracle software owner user does not exist or if you require a new Oracle software owner user, create it as follows, using the user name `oracle` unless a user with that name already exists.

1. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example, `oinstall`
 - The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba ,oper`
2. Set the password of the `oracle` user:

```
# passwd oracle
```

See the "[Verifying That the UNIX User nobody Exists](#)" section on page 5-12 to continue.

Modify an Existing Oracle Software Owner User

If the `oracle` user exists, but its primary group is not `oinstall` or it is not a member of the appropriate OSDBA or OSOPER groups, you can modify it by entering a command similar to the following, specifying the primary group using the `-g` option and any required secondary groups using the `-G` option:

```
# /usr/sbin/usermod -g oinstall -G dba[,oper] oracle
```

Verifying That the UNIX User nobody Exists

Before installing the software, verify that the UNIX user `nobody` exists on the system:

1. To determine whether the user exists, enter the following command:

```
# id nobody
```

If this command displays information about the `nobody` user, you do not have to create that user.

2. If the `nobody` user does not exist, enter the following command to create it:

```
# /usr/sbin/useradd nobody
```

3. Repeat this procedure on all of the other cluster nodes.

Create Identical Users and Groups on Other Cluster Nodes

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

The Oracle software owner user and the Oracle Inventory, OSDBA, and OSOPER groups must exist and be identical on all cluster nodes. To create these identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, then create the user and groups with the same name and ID on the other cluster nodes.

Identifying the User and Group IDs

To determine the user ID (UID) of the Oracle software owner user and the group IDs (GID) of the Oracle Inventory, OSDBA, and OSOPER groups, follow these steps:

1. Enter the following command:

```
# id oracle
```

The output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

2. From the output, identify the UID for the `oracle` user and the GIDs for the groups to which it belongs.

Create the User and Groups on the Other Cluster Nodes

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

1. Log in to the next cluster node as `root`.
2. Enter commands similar to the following to create the `oinstall` and `dba` groups, and if required, the `oper` group. Use the `-g` option to specify the correct GID for each group.

```
# /usr/sbin/groupadd -g 200 oinstall
# /usr/sbin/groupadd -g 201 dba
# /usr/sbin/groupadd -g 202 oper
```

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

3. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -u 200 -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-u` option specifies the user ID, which must be the user ID that you identified in the previous subsection.
- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`
- The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba,oper`

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

4. Set the password of the `oracle` user:

```
# passwd oracle
```

Configure SSH on All Cluster Nodes

Before you install and use Oracle Real Application clusters, you must configure secure shell (SSH) for the `oracle` user on all cluster nodes. The Installer uses the `ssh` and `scp` commands during installation to run remote commands on and copy files to the other cluster nodes. You must configure SSH so that these commands do not prompt for a password.

Note: This section describes how to configure OpenSSH version 3. If SSH is not available, the Installer attempts to use `rsh` and `rcp` instead. However, these services are disabled by default on most Linux systems.

To configure SSH, complete the following steps on each cluster node:

1. Log in as the `oracle` user.
2. If necessary, create the `.ssh` directory in the `oracle` user's home directory and set the correct permissions on it:

```
$ mkdir ~/.ssh
$ chmod 755 ~/.ssh
```

3. Enter the following commands to generate an RSA key for version 2 of the SSH protocol:

```
$ /usr/bin/ssh-keygen -t rsa
```

At the prompts:

- Accept the default location for the key file
- Enter and confirm a pass phrase that is different from the `oracle` user's password

This command writes the public key to the `~/.ssh/id_rsa.pub` file and the private key to the `~/.ssh/id_rsa` file. Never distribute the private key to anyone.

4. Enter the following commands to generate a DSA key for version 2 of the SSH protocol:

```
$ /usr/bin/ssh-keygen -t dsa
```

At the prompts:

- Accept the default location for the key file
- Enter and confirm a pass phrase that is different from the `oracle` user's password

This command writes the public key to the `~/.ssh/id_dsa.pub` file and the private key to the `~/.ssh/id_dsa` file. Never distribute the private key to anyone.

5. Copy the *contents* of the `~/.ssh/id_rsa.pub` and `~/.ssh/id_dsa.pub` files to the `~/.ssh/authorized_keys` file on this node and to the same file on all other cluster nodes.

Note: The `~/.ssh/authorized_keys` file on every node must contain the contents from all of the `~/.ssh/id_rsa.pub` and `~/.ssh/id_dsa.pub` files that you generated on all cluster nodes.

6. Change the permissions on the `~/.ssh/authorized_keys` file on all cluster nodes:

```
$ chmod 644 ~/.ssh/authorized_keys
```

At this point, if you use `ssh` to log in to or run a command on another node, you are prompted for the pass phrase that you specified when you created the DSA key.

To enable the Installer to use the `ssh` and `scp` commands without being prompted for a pass phrase, follow these steps:

1. On the system where you want to install the software, log in as the `oracle` user.
2. Enter the following commands:

```
$ exec /usr/bin/ssh-agent $SHELL
$ /usr/bin/ssh-add
```

3. At the prompts, enter the pass phrase for each key that you generated.

If you have configured SSH correctly, you can now use the `ssh` or `scp` commands without being prompted for a password or a pass phrase.

4. To test the SSH configuration, enter the following commands from the same terminal session, testing the configuration of each cluster node:

```
$ ssh nodename1 date
$ ssh nodename2 date
.
.
```

These commands should display the date set on each node. If any node prompts for a password or pass phrase, verify that the `~/.ssh/authorized_keys` file on that node contains the correct public keys.

Note: The first time you use SSH to connect to a node from a particular system, you might see a message stating that the authenticity of the host could not be established. Enter *yes* at the prompt to continue. You should not see this message again when you connect from this system to that node.

5. You must run the Installer from this session or remember to repeat steps 2 and 3 before you start the Installer from a different terminal session.

Configure Kernel Parameters and Shell Limits

Notes: The kernel parameter and shell limit values shown in the following sections are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

You must set the kernel parameters and shell limits on all cluster nodes.

See the following sections for information about configuring the system kernel parameters and shell limits.

- ["Configuring Kernel Parameters"](#) on page 5-16
- ["Set Shell Limits for the oracle User"](#) on page 5-18

Configuring Kernel Parameters

Verify that the kernel parameters shown in the following table are set to values greater than or equal to the recommended value shown. The procedure following the table describes how to verify and set the values.

Parameter	Value	File
semmsl	250	<code>/proc/sys/kernel/sem</code>
semms	32000	
semopm	100	
semmni	128	
shmall	2097152	<code>/proc/sys/kernel/shmall</code>
shmmax	Half the size of physical memory	<code>/proc/sys/kernel/shmmax</code>
shmmni	4096	<code>/proc/sys/kernel/shmmni</code>
file-max	65536	<code>/proc/sys/fs/file-max</code>
ip_local_port_range	1024 65000	<code>/proc/sys/net/ipv4/ip_local_port_range</code>

Note: If the current value for any parameter is higher than the value listed in this table, do not change the value of that parameter.

To view the current value specified for these kernel parameters, and to change them if necessary, follow these steps:

1. Enter the commands shown in the following to view the current values of the kernel parameters:

Note: Make a note of the current values and identify any values that you must change.

Parameter	Command
semmsl, semmns, semopm, and semmni	# /sbin/sysctl -a grep sem This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and shmmni	# /sbin/sysctl -a grep shm
file-max	# /sbin/sysctl -a grep file-max
ip_local_port_range	# /sbin/sysctl -a grep ip_local_port_range This command displays a range of port numbers.

2. If the value of any kernel parameter is different to the recommended value, complete the following steps:
 - a. Using any text editor, create or edit the `/etc/sysctl.conf` file and add or edit lines similar to the following:

Note: Include lines only for the kernel parameter values that you want to change. For the semaphore parameters (`kernel.sem`), you must specify all four values. However, if any of the current values are larger than the recommended value, specify the larger value.

```
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.file-max = 65536
net.ipv4.ip_local_port_range = 1024 65000
```

By specifying the values in the `/etc/sysctl.conf` file, they persist when you reboot the system.

- b. Enter the following command to change the current values of the kernel parameters:

```
# /sbin/sysctl -p
```

Review the output from this command to verify that the values are correct. If the values are incorrect, edit the `/etc/sysctl.conf` file, then enter this command again.

- c. On UnitedLinux only, enter the following command to cause the system to read the `/etc/sysctl.conf` file when it reboots:

```
# /sbin/chkconfig boot.sysctl on
```

3. Repeat this procedure on all other cluster nodes.

Set Shell Limits for the oracle User

To improve the performance of the software, you must increase the following shell limits for the `oracle` user:

Shell Limit	Item in <code>limits.conf</code>	Hard Limit
Maximum number of open file descriptors	<code>nofile</code>	65536
Maximum number of processes available to a single user	<code>nproc</code>	16384

To increase the shell limits:

1. Add the following lines to `/etc/security/limits.conf` file:

```
*          soft   nproc    2047
*          hard   nproc    16384
*          soft   nofile   1024
*          hard   nofile   65536
```

2. Add the following line to the `/etc/pam.d/login` file, if it does not already exist:

```
session    required    /lib/security/pam_limits.so
```

3. Depending on the `oracle` user's default shell, make the following changes to the default shell start-up file:

- For the Bourne, Bash, or Korn shell, add the following lines to the `/etc/profile` file (or the `/etc/profile.local` file on UnitedLinux systems):

```
if [ $USER = "oracle" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
        ulimit -p 16384
        ulimit -n 65536
    else
        ulimit -u 16384 -n 65536
    fi
fi
```

- For the C or tcsh shell, add the following lines to the `/etc/csh.login` file (or the `/etc/csh.login.local` file on UnitedLinux systems):

```
if ( $USER == "oracle" ) then
    limit maxproc 16384
    limit descriptors 65536
endif
```

4. Repeat this procedure on all other cluster nodes.

Identify Required Software Directories

You must identify or create four directories for the Oracle software, as follows:

- Oracle base directory
- Oracle Inventory directory
- CRS home directory
- Oracle home directory

The following subsections describe the requirements for these directories.

Oracle Base Directory

The Oracle base directory acts as a top-level directory for Oracle software installations. It is analogous to the `C:\Oracle` directory used for Oracle software installations on Windows systems. On UNIX systems, the Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to the following for the Oracle base directory:

```
/mount_point/app/oracle_sw_owner
```

In this example:

- *mount_point* is the mount point directory for the file system that will contain the Oracle software.

The examples in this guide use `/u01` for the mount point directory. However, you could choose another mount point directory, `/oracle` or `/opt/oracle` for example.

- *oracle_sw_owner* is the UNIX user name of the Oracle software owner, for example `oracle`.

You can use the same Oracle base directory for more than one installation or you can create separate Oracle base directories for different installations. If different UNIX users install Oracle software on the same system, each user must create a separate Oracle base directory. The following example Oracle base directories could all exist on the same system:

```
/u01/app/oracle  
/u01/app/orauser  
/opt/oracle/app/oracle
```

The following sections describe how to identify existing Oracle base directories that might be suitable for your installation and how to create a new Oracle base directory if necessary.

Regardless of whether you create a new Oracle base directory or decide to use an existing one, you must set the `ORACLE_BASE` environment variable to specify the full path to this directory.

Note: The Oracle base directory can be on a local file system or on an NFS file system on a certified NAS device. Do not create the Oracle base directory on an OCFS version 1 file system.

Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system. It is required by, and shared by, all Oracle software installations on a single system. The first time you install Oracle software on a system, the Installer prompts you to specify the path to this directory. If you are installing the software on a local file system, Oracle recommends that you choose the following path:

```
oracle_base/oraInventory
```

The Installer creates the directory that you specify and sets the correct owner, group, and permissions on it. You do not need to create it.

Note: All Oracle software installations rely on this directory. Make sure that you back it up regularly.

Do not delete this directory unless you have completely removed all Oracle software from the system.

CRS Home Directory

The CRS home directory is the directory where you choose to install the software for Oracle Cluster Ready Services. You must install CRS in a separate home directory. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
oracle_base/product/10.1.0/crs
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Oracle Home Directory

The Oracle home directory is the directory where you choose to install the software for a particular Oracle product. You must install different Oracle products, or different releases of the same Oracle product, in separate Oracle home directories. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the Oracle home directory:

```
oracle_base/product/10.1.0/db_1
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Identify or Create an Oracle Base Directory

Before starting the installation, you must either identify an existing Oracle base directory or if required, create a new one. This section contains information about the following:

- [Identifying an Existing Oracle Base Directory](#)
- [Creating a New Oracle Base Directory](#)

Note: You can choose to create a new Oracle base directory, even if other Oracle base directories exist on the system.

Identifying an Existing Oracle Base Directory

Existing Oracle base directories might not have paths that comply with OFA guidelines. However, if you identify an existing Oracle Inventory directory or existing Oracle home directories, you can usually identify the Oracle base directories, as follows:

- Identify an existing Oracle Inventory directory

Enter the following command on all cluster nodes to view the contents of the `oraInst.loc` file:

```
# more /etc/oraInst.loc
```

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inventory_loc` parameter identifies the Oracle Inventory directory (`oraInventory`) on that system. The parent directory of the `oraInventory` directory is typically an Oracle base directory. In the previous example, `/u01/app/oracle` is an Oracle base directory.

- Identify existing Oracle home directories

Enter the following command on all cluster nodes to view the contents of the `oratab` file:

```
# more /etc/oratab
```

If the `oratab` file exists, it contains lines similar to the following:

```
*:/u03/app/oracle/product/10.1.0/db_1:N
*/opt/orauser/infra_904:N
*/oracle/9.2.0:N
```

The directory paths specified on each line identify Oracle home directories. Directory paths that end with the user name of the Oracle software owner that you want to use are valid choices for an Oracle base directory. If you intend to use the `oracle` user to install the software, you could choose one of the following directories from the previous example:

```
/u03/app/oracle
/oracle
```

Note: If possible, choose a directory path similar to the first (`/u03/app/oracle`). This path complies with the OFA guidelines.

Before deciding to use an existing Oracle base directory for this installation, make sure that it satisfies the following conditions:

- It should not be on the same file system as the operating system.
- It must have an identical path on all cluster nodes or it must be an NFS file system on a certified NAS device.

If you are not using an NFS file system, create identical Oracle base directories on the other nodes.

- It must have sufficient free disk space on all cluster nodes, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB

Use the following command to determine the free disk space on the file system where the Oracle base directory is located:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the ["Creating a New Oracle Base Directory"](#) section on page 5-22.
- If an Oracle base directory exists and you want to use it, see the ["Configure Disk Storage for Oracle CRS and Database Files"](#) section on page 5-23.

When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify the directory you chose.

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the following section.

Creating a New Oracle Base Directory

Before you create a new Oracle base directory, you must identify an appropriate file system with sufficient free disk space, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB

To identify an appropriate file system, follow these steps:

1. Use the `df -k` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has appropriate free space.

Note: The file system can be a local file system or an NFS file system on a certified NAS device. Do not create the Oracle base directory on an OCFS file system.

The path to the Oracle base directory must be the same on all nodes.

3. Note the name of the mount point directory for the file system that you identified.

To create the Oracle base directory and specify the correct owner, group, and permissions for it, follow these steps:

1. Enter commands similar to the following to create the recommended subdirectories in the mount point directory that you identified and set the appropriate owner, group, and permissions on them:

```
# mkdir -p /mount_point/app/oracle_sw_owner
# chown -R oracle:oinstall /mount_point/app
# chmod -R 775 /mount_point/app
```

If the mount point you identified is `/u01` and `oracle` is the user name of the Oracle software owner, the recommended Oracle base directory path is as follows:

```
/u01/app/oracle
```

2. If necessary, repeat the commands listed in the previous step to create the same directory on the other cluster nodes.
3. When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify this directory.

Configure Disk Storage for Oracle CRS and Database Files

This section describes the storage options available to you for storing Oracle Cluster Ready Services files, Oracle Database Files, and optionally, Oracle Database recovery files. After you choose the storage method that you want to use for each file type, see the following sections for information about configuring the required storage.

Note: You do not have to use the same storage option for each type of file.

Choosing a Storage Option for Oracle CRS Files

Before you install Oracle Cluster Ready Services, you must choose the storage option that you want to use for the Oracle Cluster Registry (100 MB) and CRS voting disk (20 MB). You cannot use Automatic Storage Management to store these files, because they must be accessible before any Oracle instance starts.

Apart from ASM, you can use any of the storage options listed in the next section for the Oracle CRS files.

Choosing a Storage Option for Oracle Database Files

If you want to create a database during the installation, you must choose one of the following storage options for the database files:

- File system
- Automatic Storage Management
- Raw Devices
- Cluster file system using OCFS for Linux
- Automatic Storage Management
- NFS file system on a certified NAS device
- Raw partitions

Note: For the most up-to-date information about supported storage options for RAC installations, see the Certify pages on the *OracleMetaLink* Web site:

<http://metalink.oracle.com>

Choosing a Storage Option for Oracle Database Recovery Files

Additionally, if you want to enable automated backups during the installation, you must choose one of the following storage options for recovery files (the flash recovery area):

- Cluster file system using OCFS for Linux
- Automatic Storage Management
- NFS file system on a certified NAS device

The storage option that you choose for recovery files can be the same as or different to the option you choose for the database files.

Configuring Disk Storage

For information about how to configure disk storage before you start the installation, see one of the following sections depending on your choice:

- To use a file system for Oracle CRS, database, or recovery file storage, see the ["Create Directories for Oracle CRS, Database, or Recovery Files"](#) section on page 5-24.
- To use ASM for database or recovery file storage, see the ["Configure Disks for Automatic Storage Management"](#) section on page 5-27.
- To use raw devices (partitions or logical volumes) for Oracle CRS or database file storage, see the ["Configure Raw Partitions"](#) section on page 5-36.

Create Directories for Oracle CRS, Database, or Recovery Files

If you decide to place the Oracle CRS, database, or recovery files on a file system, use the following guidelines when deciding where to place them.

Guidelines for Placing Oracle CRS Files on a File System

The Installer does not suggest a default location for the Oracle Cluster Registry (OCR) or the Oracle CRS voting disk. If you choose to create these files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.
- It must have at least 100 MB of free disk space for the OCR and 20 MB of free disk space for the CRS voting disk.
- For improved reliability, you should choose a file system on a highly available storage device, for example, a RAID device that implements mirroring.
- If you are placing the Oracle Cluster Ready Services software on a shared file system, you can use the same file system for these files.
- The `oracle` user must have write permissions to create the files in the path you specify.

Note: If you are upgrading from Oracle9i Release 2, you can continue to use the raw device or shared file that you used for the SRVM configuration repository instead of creating a new file for the OCR.

Guidelines for Placing Oracle Database Files on a File System

If you choose to place the Oracle database files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.
- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.
This default location is not recommended for production databases.
- You can choose either a single file system or more than one file system to store the database files:
 - If you want to use a single file system, choose a file system on a physical device that is dedicated to the database.
For best performance and reliability, choose a RAID device or a logical volume on more than one physical device and implement the stripe-and-mirror-everything (SAME) methodology.
 - If you want to use more than one file system, choose file systems on separate physical devices that are dedicated to the database.
This method enables you to distribute physical I/O and create separate control files on different devices for increased reliability. It also enables you to fully implement the OFA guidelines. You must choose either the Advanced database creation option or the Custom installation type during the installation to implement this method.
- If you intend to create a preconfigured database during the installation, the file system (or file systems) that you choose must have at least 1.2 GB of free disk space.
For production databases, you must estimate the disk space requirement depending on the use you want to make of the database.
- For optimum performance, the file systems that you choose should be on physical devices that are used only by the database.
- The `oracle` user must have write permissions to create the files in the path you specify.

Guidelines for Placing Oracle Recovery Files on a File System

Note: You must choose a location for recovery files only if you intend to enable automated backups during the installation.

If you choose to place the Oracle recovery files on a file system, use the following guidelines when deciding where to place them:

- To prevent disk failure from making both the database files and the recovery files unavailable, place the recovery files in a file system on a different physical disk from the database files.

Note: Alternatively use an ASM disk group with a normal or high redundancy level for either or both file types.

- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.

- The file system that you choose should have at least 2 GB of free disk space. The disk space requirement is the default disk quota configured for the flash recovery area (specified by the `DB_RECOVERY_FILE_DEST_SIZE` initialization parameter).

If you choose the Custom installation type or the Advanced database configuration option, you can specify a different disk quota value. After you create the database, you can also use Oracle Enterprise Manager Grid Control or Database Control to specify a different value.

For more information about sizing the flash recovery area, see the *Oracle Backup and Recovery Basics* manual.
- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.

This default location is not recommended for production databases.
- The `oracle` user must have write permissions to create the files in the path you specify.

Creating Required Directories

Note: You must complete this procedure only if you want to place the Oracle CRS, database, or recovery files on a separate file system to the Oracle base directory.

To create directories for the Oracle CRS, database, or recovery files on separate file systems to the Oracle base directory, follow these steps:

1. If necessary, configure the shared file systems that you want to use and mount them on each node.

Note: Make sure that the file systems are configured to mount automatically when a node reboots.

2. Use the `df -k` command to determine the free disk space on each mounted file system.
3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
CRS files	Choose a file system with at least 120 MB of free disk space
Datafiles	Choose either: <ul style="list-style-type: none"> ■ A single file system with at least 1.2 GB of free disk space ■ Two or more file systems with at least 1.2 GB of free disk space in total
Recovery files	Choose a file system with at least 2 GB of free disk space.

If you are using the same file system for more than one type of file, add the disk space requirements for each type to determine the total disk space requirement.

4. Note the names of the mount point directories for the file systems that you identified.
5. Enter commands similar to the following to create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:
 - CRS file directory:


```
# mkdir /mount_point/oracrs
# chown oracle:oinstall /mount_point/oracrs
# chmod 775 /mount_point/oracrs
```
 - Database file directory:


```
# mkdir /mount_point/oradata
# chown oracle:oinstall /mount_point/oradata
# chmod 775 /mount_point/oradata
```
 - Recovery file directory (flash recovery area):


```
# mkdir /mount_point/flash_recovery_area
# chown oracle:oinstall /mount_point/flash_recovery_area
# chmod 775 /mount_point/flash_recovery_area
```
6. If you also want to use ASM or raw devices for storage, see one of the following sections:
 - [Configure Disks for Automatic Storage Management](#)
 - [Configure Raw Partitions](#)

Otherwise see the "[Verify that the Required Software is Running](#)" section on page 5-41.

Configure Disks for Automatic Storage Management

This section describes how to configure disks for use with ASM. Before you configure the disks, you must determine the number of disks and the amount of free disk space that you require. The following sections describe how to identify the requirements and configure the disks:

- [Identifying Storage Requirements for ASM](#)
- [Using an Existing ASM Disk Group](#)
- [Configuring Disks for ASM](#)

Note: Although this section refers to disks, you can also use zero-padded files on a certified NAS storage device in an ASM disk group. See the appendices in *Oracle Database Installation Guide for UNIX Systems* for information about creating and configuring NAS-based files for use in an ASM disk group.

Identifying Storage Requirements for ASM

To identify the storage requirements for using ASM, you must determine how many devices and the amount of free disk space that you require. To complete this task, follow these steps:

1. Determine whether you want to use ASM for Oracle database files (datafiles), recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. One can use the file system, while the other uses ASM.

If you choose to enable automated backups and you do not have a cluster file system available, you must choose ASM for recovery file storage.

If you enable automated backups during the installation, you can choose ASM as the storage mechanism for recovery files by specifying an ASM disk group for the flash recovery area. Depending how you choose to create a database during the installation, you have the following options:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to use the same ASM disk group for datafiles and recovery files, or you can choose to use different disk groups for each file type.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must use the same ASM disk group for database files and recovery files.

2. Choose the ASM redundancy level that you want to use for the ASM disk group.

The redundancy level that you choose for the ASM disk group determines how ASM mirrors files in the disk group and determines the number of disks and amount of disk space that you require, as follows:

- 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 ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use only RAID or similar devices that provide their own data protection mechanisms as disk devices in this type of disk group.

- Normal redundancy

In a normal redundancy disk group, ASM uses two-way mirroring by default, to increase performance and reliability. 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 most installations, Oracle recommends that you use normal redundancy disk groups.

- High redundancy

In a high redundancy disk group, 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.

While high redundancy disk groups do provide a high level of data protection, you must consider the higher cost of additional storage devices before deciding to use this redundancy level.

3. Determine the total amount of disk space that you require for the datafiles and recovery files.

Use the following table to determine the minimum number of disks and the minimum disk space requirements for the installation:

Redundancy Level	Minimum Number of Disks	Datafiles	Recovery Files	Both File Types
External	1	1.15 GB	2.3 GB	3.45 GB
Normal	2	2.3 GB	4.6 GB	6.9 GB
High	3	3.45 GB	6.9 GB	10.35 GB

You must also add additional disk space for the ASM metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

$$15 + (2 * \text{number_of_disks}) + (126 * \text{number_of_ASM_instances})$$

For example, for a four-node RAC installation, using three disks in a high redundancy disk group, you require an additional 525 MB of disk space:

$$(15 + (2 * 3) + (126 * 4)) = 525$$

If an existing ASM instance exists on the system, 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 installation.

The following section describes how to identify existing disk groups and determine the free disk space that they contain.

4. Optionally identify failure groups for the ASM disk group devices.

Note: You need to complete this step only if you intend to use an installation method that runs DBCA in interactive mode, for example, if you intend to choose the Custom installation type or the Advanced database configuration option. Other installation types do not enable you to specify failure groups.

If you intend to use a normal or high redundancy disk group, 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 SCSI controller, the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To avoid failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

5. If you are sure that a suitable disk group does not exist on the system, 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 ASM disk group should be the same size and have the same performance characteristics.
 - Do not specify more than one partition on a single physical disk as a disk group device. 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 ASM disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing ASM from optimizing I/O across the physical devices.

For information about completing this task, see the "[Configure Disks for Automatic Storage Management](#)" section on page 5-27.

Using an Existing ASM Disk Group

If you want to use ASM as the storage option for either database or recovery files, and an existing ASM disk group exists, you have the following choices, depending on the installation method that you select:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to create a new disk group or use an existing one.

The same choice is available to you if you use DBCA after the installation to create a database.
- If you select an installation type that runs DBCA in non-interactive mode, you must choose an existing disk group for the new database; you cannot create a new disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The ASM instance that manages the existing disk group can be running in a different Oracle home directory.

To determine whether an existing ASM disk group exists, or to determine whether there is sufficient disk space in a disk group, you can use Oracle Enterprise Manager Database Grid Control or Database Control. Alternatively, you can use the following procedure:

1. View the contents of the `oratab` file to determine whether an ASM instance is configured on the system:

```
# more /etc/oratab
```

If an ASM instance is configured on the system, the `oratab` file should contain a line similar to the following:

```
+ASM:oracle_home_path:N
```

In this example, `+ASM` is the system identifier (SID) of the ASM instance and `oracle_home_path` is the Oracle home directory where it is installed. By convention, the SID for an ASM instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the ASM instance that you want to use.
3. Connect to the ASM instance as the `SYS` user with `SYSDBA` privilege and start the instance if necessary:

```
# $ORACLE_HOME/bin/sqlplus "SYS/SYS_password as SYSDBA"
SQL> STARTUP
```

4. Enter the following command to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:


```
SQL> SELECT NAME,TYPE,TOTAL_MB,FREE_MB FROM V$ASM_DISKGROUP;
```
5. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.
6. If necessary, install or identify the additional disk devices required to meet the storage requirements listed in the previous section.

Note: If you are adding devices to an existing disk group, Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Configuring Disks for ASM

Oracle provides an ASM library driver that you can use to improve the I/O performance of databases that use ASM for storage management. If you intend to use ASM for database storage, Oracle recommends that you install the ASM library driver and associated utilities and use them to configure the devices that you want to include in an ASM disk group. If you do not use the ASM library driver, you must bind each disk device that you want to use to a raw device. The following sections describe how to configure disks using both methods:

- [Configuring Disks for ASM Using the ASM Library Driver](#)
- [Configuring Disks for ASM Using Raw Devices](#)

Configuring Disks for ASM Using the ASM Library Driver

To use the ASM library driver to configure ASM devices, complete the following tasks.

Install and Configure the ASM Library Driver Software

To install and configure the ASM library driver, and utilities, follow these steps:

1. Switch user to the `root` user:

```
# su -
```

2. Enter the following command to determine the kernel version and architecture of the system:

```
# uname -rm
```

3. If you are not installing the software from a hard disk, insert the Oracle Database 10g Companion CD-ROM or the Oracle Database 10g DVD-ROM into the drive.

If the disc does not mount automatically, see the "Mounting Discs on Linux" section in the appendices of *Oracle Database Installation Guide for UNIX Systems*.

4. Change directory to the directory containing the ASM library driver packages for your Linux distribution and architecture:

```
$ cd directory_path/RPMS/distribution
```

In this example, *distribution* is the distribution of Linux that you are using and *directory_path* is one of the following:

- The CD-ROM mount point directory path
 - The Oracle Database 10g Companion CD installation directory (*companion*) on the DVD-ROM
 - The path of the `disk1` directory on the hard disk
5. Enter commands similar to the following to install the packages:

```
# rpm -Uvh oracleasm-support-1.0.0-1.arch.rpm oracleasm-lib-1.0.0-1.arch.rpm
# rpm -Uvh oracleasm-kernel-version-1.0.0-1.arch.rpm
```

In this example, *arch* is the system architecture type and *kernel_version* is the version of the kernel that you are using. For example, if you are using the Red Hat Enterprise Linux AS 2.1 enterprise kernel on a system with the i686 architecture, enter commands similar to the following:

```
# rpm -Uvh oracleasm-support-1.0.0-1.i686.rpm oracleasm-lib-1.0.0-1.i686.rpm
# rpm -Uvh oracleasm-2.4.9-e-enterprise-1.0.0-1.i686.rpm
```

6. Enter a command similar to the following to determine the UID of the Oracle software owner user that you are using for this installation (typically `oracle`) and the GID of the OSDBA group (typically `dba`):

```
# id oracle
```

7. Enter the following command to run the `oracleasm` initialization script with the `configure` option:

```
# /etc/init.d/oracleasm configure
```

8. Enter the following information in response to the prompts that the script displays:

Prompt	Suggested Response
Default UID to own the driver interface:	Specify the UID of the Oracle software owner user (<code>oracle</code>).
Default GID to own the driver interface:	Specify the GID of the OSDBA group (<code>dba</code>).
Start Oracle ASM Library driver on boot (y/n):	Enter <code>y</code> to start the Oracle ASM library driver when the system starts.

The script completes the following tasks:

- Creates the `/etc/sysconfig/oracleasm` configuration file
- Creates the `/dev/oracleasm` mount point
- Loads the `oracleasm` kernel module

- Mounts the ASM library driver file system

Note: The ASM library driver file system is not a regular file system. It is used only by the ASM library to communicate with the ASM driver.

9. Repeat this procedure on all cluster nodes where you want to install RAC.

Configure the Disk Devices to Use the ASM Library Driver

To configure the disk devices that you want to use in an ASM disk group, complete the following steps:

1. If necessary, install the shared disks that you intend to use for the disk group and reboot the system.
2. To identify the device name for the disks that you want to use, enter the following command:

```
# /sbin/fdisk -l
```

Depending on the type of disk, the device name can vary:

Disk Type	Device Name Format	Description
IDE disk	/dev/hdxn	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, /dev/hda is the first disk on the first IDE bus.
SCSI disk	/dev/sdxn	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, /dev/sda is the first disk on the first SCSI bus.
RAID disk	/dev/rd/cxdypz /dev/ida/cxdypz	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <i>x</i> is a number that identifies the controller, <i>y</i> is a number that identifies the disk and <i>z</i> is a number that identifies the partition. For example, /dev/ida/c0d1 is the second logical drive on the first controller.

To include devices in a disk group, you can specify either whole drive device names or partition device names.

Note: Oracle recommends that you create a single whole-disk partition on each disk that you want to use.

3. Using either `fdisk` or `parted`, create a single whole-disk partition on the disk device that you want to use.
4. Enter a command similar to the following to mark a disk as an ASM disk:

```
/etc/init.d/oracleasm createdisk disk1 /dev/sdb1
```

In this example, `disk1` is a name that you want to assign to the disk.

Note: If you are using a multi-pathing disk driver with ASM, make sure that you specify the correct logical device name for the disk.

- To make the disk available on the other cluster nodes, enter the following command as root on each node:

```
# /etc/init.d/oracleasm scandisks
```

This command identifies shared disks attached to the node that are marked as ASM disks.

Administer the ASM Library Driver and Disks

To administer the ASM library driver and disks, use the `oracleasm` initialization script with different options, as follows:

Option	Description
configure	Use the <code>configure</code> option to reconfigure the ASM library driver, if necessary: <pre># /etc/init.d/oracleasm configure</pre>
enable disable	Use the <code>disable</code> and <code>enable</code> options to change the behavior of the ASM library driver when the system boots. The <code>enable</code> option causes the ASM library driver to load when the system boots: <pre># /etc/init.d/oracleasm enable</pre>
start stop restart	Use the <code>start</code> , <code>stop</code> , and <code>restart</code> options to load or unload the ASM library driver without rebooting the system: <pre># /etc/init.d/oracleasm restart</pre>
createdisk	Use the <code>createdisk</code> option to mark a disk device for use with the ASM library driver and give it a name: <pre># /etc/init.d/oracleasm createdisk diskname devicename</pre>
deletedisk	Use the <code>deletedisk</code> option to unmark a named disk device: <pre># /etc/init.d/oracleasm deletedisk diskname</pre> Caution: Do not use this command to unmark disks that are being used by an ASM disk group. You must drop the disk from the ASM disk group before you unmark it.
querydisk	Use the <code>querydisk</code> option to determine whether a disk device or disk name is being used by the ASM library driver: <pre># /etc/init.d/oracleasm querydisk {diskname devicename}</pre>
listdisks	Use the <code>listdisks</code> option to list the disk names of marked ASM library driver disks: <pre># /etc/init.d/oracleasm listdisks</pre>
scandisks	Use the <code>scandisks</code> option to enable cluster nodes to identify which shared disks have been marked as ASM library driver disks on another node: <pre># /etc/init.d/oracleasm scandisks</pre>

Configuring Disks for ASM Using Raw Devices

Note: For improved performance and easier administration, Oracle recommends that you use the ASM library driver rather than raw devices to configure ASM disks.

To configure disks for ASM using raw devices, follow these steps:

- If necessary, install the shared disks that you intend to use for the disk group and reboot the system.

2. To identify the device name for the disks that you want to use, enter the following command:

```
# /sbin/fdisk -l
```

Depending on the type of disk, the device name can vary:

Disk Type	Device Name Format	Description
IDE disk	/dev/hdxn	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, /dev/hda is the first disk on the first IDE bus.
SCSI disk	/dev/sdxn	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, /dev/sda is the first disk on the first SCSI bus.
RAID disk	/dev/rd/cxdypz /dev/ida/cxdypz	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <i>x</i> is a number that identifies the controller, <i>y</i> is a number that identifies the disk and <i>z</i> is a number that identifies the partition. For example, /dev/ida/c0d1 is the second logical drive on the first controller.

To include devices in a disk group, you can specify either whole drive device names or partition device names.

Note: Oracle recommends that you create a single whole disk partition on each disk that you want to use.

3. Using either `fdisk` or `parted`, create a single whole-disk partition on the disk device that you want to use.
4. On Red Hat Linux, complete the following steps to bind the disk devices to raw devices:
 - a. To determine what raw devices are already bound to other devices, enter the following command:

```
# /usr/sbin/raw -qa
```

Raw devices have device names in the form `/dev/raw/rawn`, where *n* is a number that identifies the raw device.

- b. Open the `/etc/sysconfig/rawdevices` file in any text editor and add a line similar to the following for each device that you want to include in a disk group:

```
/dev/raw/raw1 /dev/sdb1
```

Note: If you are using a multi-pathing disk driver with ASM, ensure that specify the correct logical device name for the disk.

Specify an unused raw device for each disk device.

- c. For each raw device that you specified in the `rawdevices` file, enter commands similar to the following to set the owner, group, and permissions on the device file:

```
# chown oracle:dba /dev/raw/rawn
```

```
# chmod 660 /dev/raw/rawn
```

- d. To bind the partitions to the raw devices, enter the following command:

```
# /sbin/service rawdevices restart
```

By editing the `rawdevices` file, the system binds the partitions to the raw devices when it reboots.

5. On UnitedLinux, complete the following steps to bind the disk devices to raw devices:

- a. To determine what raw devices are already bound to other devices, enter the following command:

```
# /usr/sbin/raw -qa
```

Raw devices have device names in the form `/dev/raw/rawn`, where `n` is a number that identifies the raw device.

- b. Open the `/etc/raw` file in any text editor and add a line similar to the following for each device that you want to include in a disk group:

```
raw1:sdb1
```

Note: If you are using a multi-pathing disk driver with ASM, ensure that you specify the correct logical device name for the disk.

Specify an unused raw device for each disk device.

- c. For each raw device that you specified in the `/etc/raw` file, enter commands similar to the following to set the owner, group, and permissions on the device file:

```
# chown oracle:dba /dev/raw/rawn
# chmod 660 /dev/raw/rawn
```

- d. To bind the disk devices to the raw devices, enter the following command:

```
# /etc/init.d/raw
```

By editing the `/etc/raw` file, the system binds the partitions to the raw devices when it reboots.

- e. To ensure that the raw devices are bound when the system reboots, enter the following command:

```
# /sbin/chkconfig raw on
```

Configure Raw Partitions

See the following sections for information about configuring raw partitions:

- ["Review Important Information"](#) on page 5-37
- ["Create Raw Partitions"](#) on page 5-37
- ["Bind the Partitions to Raw Devices"](#) on page 5-39
- ["Create the DBCA Raw Device Mapping File"](#) on page 5-40

Note: If you are using ASM for database file storage, you need only create raw devices for the Oracle CRS files. However, Oracle recommends that you use that file system to store the Oracle CRS files instead of using raw devices for them.

Review Important Information

The procedures contained in this section describe how to create raw partitions for Oracle CRS and database file storage. Although Red Hat Enterprise Linux 3 and UnitedLinux 1.0 provide a Logical Volume Manager (LVM), this LVM is not cluster aware. For this reason, Oracle does not support the use of logical volumes with RAC for either CRS or database files.

Note: The use of logical volumes for raw devices is supported only for single-instance databases. They are not supported for RAC databases.

See Also: For more information about completing steps 1 and 2, see the operating system documentation or the `lvm` man page.

Create Raw Partitions

To create the required raw partitions, follow these steps:

1. If necessary, install the shared disks that you intend to use for the database and reboot the system.

Note: Because the number of partitions that you can create on a single device is limited, you might need to create the required raw partitions on more than one device.

2. To identify the device name for the disks that you want to use for the database, enter the following command:

```
# /sbin/fdisk -l
```

Depending on the type of disk, the device name can vary:

Disk Type	Device Name Format	Description
IDE disk	<code>/dev/hdxn</code>	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, <code>/dev/hda</code> is the first disk on the first IDE bus.
SCSI disk	<code>/dev/sdxn</code>	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, <code>/dev/sda</code> is the first disk on the first SCSI bus.
RAID disk	<code>/dev/rd/cxdypz</code> <code>/dev/ida/cxdypz</code>	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <i>x</i> is a number that identifies the controller, <i>y</i> is a number that identifies the disk and <i>z</i> is a number that identifies the partition. For example, <code>/dev/ida/c0d1</code> is the second logical drive on the first controller.

You can create the required raw partitions either on new devices that you added or on previously partitioned devices that have unpartitioned free space. To identify devices that have unpartitioned free space, examine the start and end cylinder numbers of the existing partitions and determine whether the device contains unused cylinders.

- To create raw partitions on a device, enter a command similar to the following:

```
# /sbin/fdisk devicename
```

Use the following guidelines when creating partitions:

- Use the `p` command to list the partition table of the device.
- Use the `n` command to create a new partition.
- After you have created all the required partitions on this device, use the `w` command to write the modified partition table to the device.
- See the `fdisk` man page for more information about creating partitions.

The following table lists the number, size, and purpose of the partitions that you must create:

Number	Partition Size (MB)	Purpose
Raw Partitions for Oracle Database		
1	500	SYSTEM tablespace
1	300 + (Number of instances * 250)	SYSAUX tablespace
Number of instances	500	UNDOTBS <i>n</i> tablespace (One tablespace for each instance)
1	160	EXAMPLE tablespace
1	120	USERS tablespace
2 * number of instances	120	Two online redo log files for each instance
2	110	First and second control files
1	250	TEMP tablespace
1	5	Server parameter file (SPFILE)
1	5	Password file
Raw Logical Volumes for Oracle Cluster Ready Services (CRS)		
1	100	Oracle Cluster Registry Note: You need to create this raw partition only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle cluster registry. If you are upgrading from Oracle9i Release 2, you can continue to use the raw device that you used for the SRVM configuration repository instead of creating this new raw device.
1	20	Oracle CRS voting disk Note: You need to create this raw partition only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle CRS voting disk.

Note: If you prefer to use manual, instead of automatic, undo management, create a single RBS tablespace raw device at least 500 MB in size instead of the UNDOTBS*n* raw devices.

Bind the Partitions to Raw Devices

After you have created the required partitions, you must bind the partitions to raw devices. However, you must first determine what raw devices are already bound to other devices. The procedure that you must follow to complete this task varies, depending on the distribution that you are using:

- Red Hat:

1. To determine what raw devices are already bound to other devices, enter the following command:

```
# /usr/sbin/raw -qa
```

Raw devices have device names in the form `/dev/raw/rawn`, where *n* is a number that identifies the raw device.

2. Open the `/etc/sysconfig/rawdevices` file in any text editor and add a line similar to the following for each partition that you created:

```
/dev/raw/raw1 /dev/sdb1
```

Specify an unused raw device for each partition.

3. For the raw device that you created for the Oracle Cluster Registry, enter commands similar to the following to set the owner, group, and permissions on the device file:

```
# chown root:dba /dev/dev/rawn
# chmod 640 /dev/raw/rawn
```

4. For each other raw device that you specified in the `rawdevices` file, enter commands similar to the following to set the owner, group, and permissions on the device file:

```
# chown oracle:oinstall /dev/dev/rawn
# chmod 660 /dev/raw/rawn
```

5. To bind the partitions to the raw devices, enter the following command:

```
# /sbin/service rawdevices restart
```

By editing the `rawdevices` file, the system binds the partitions to the raw devices when it reboots.

- UnitedLinux:

1. To determine what raw devices are already bound to other devices, enter the following command:

```
# /usr/sbin/raw -qa
```

Raw devices have device names in the form `/dev/raw/rawn`, where *n* is a number that identifies the raw device.

2. Open the `/etc/raw` file in any text editor and add a line similar to the following to associate each partition with an unused raw device:

```
raw1:sdb1
```

3. For the raw device that you created for the Oracle Cluster Registry, enter commands similar to the following to set the owner, group, and permissions on the device file:

```
# chown root:dba /dev/dev/rawn
# chmod 640 /dev/raw/rawn
```

4. For each other raw device that you specified in the `/etc/raw` file, enter commands similar to the following to set the owner, group, and permissions on the device file:

```
# chown oracle:dba /dev/raw/rawn
# chmod 660 /dev/raw/rawn
```

5. To bind the partitions to the raw devices, enter the following command:

```
# /etc/init.d/raw
```

6. To ensure that the raw devices are bound when the system reboots, enter the following command:

```
# /sbin/chkconfig raw on
```

Create the DBCA Raw Device Mapping File

Note: You must complete this procedure only if you are using raw devices for database files. You do not specify the raw devices for the Oracle CRS files in the DBCA raw device mapping file.

To enable Database Configuration Assistant (DBCA) to identify the appropriate raw partition for each database file, you must create a raw device mapping file, as follows:

1. Set the `ORACLE_BASE` environment variable to specify the Oracle base directory that you identified or created previously:

- Bourne, Bash or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle ; export ORACLE_BASE
```

- C shell:

```
$ setenv ORACLE_BASE /u01/app/oracle
```

2. Create a database file subdirectory under the Oracle base directory and set the appropriate owner, group, and permissions on it:

```
# mkdir -p $ORACLE_BASE/oradata/dbname
# chown -R oracle:oinstall $ORACLE_BASE/oradata
# chmod -R 775 $ORACLE_BASE/oradata
```

In this example, `dbname` is the name of the database that you chose previously.

3. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
4. Edit the `dbname_raw.conf` file in any text editor to create a file similar to the following:

Note: The following example shows a sample mapping file for a two-instance RAC cluster.

```

system=/dev/raw/raw1
sysaux=/dev/raw/raw2
example=/dev/raw/raw3
users=/dev/raw/raw4
temp=/dev/raw/raw5
undotbs1=/dev/raw/raw6
undotbs2=/dev/raw/raw7
redo1_1=/dev/raw/raw8
redo1_2=/dev/raw/raw9
redo2_1=/dev/raw/raw10
redo2_2=/dev/raw/raw11
control1=/dev/raw/raw12
control2=/dev/raw/raw13
spfile=/dev/raw/raw14
pwdfile=/dev/raw/raw15

```

Use the following guidelines when creating or editing this file:

- Each line in the file must have the following format:

```
database_object_identifier=raw_device_path
```

- For a single instance database, the file must specify one automatic undo tablespace datafile (`undotbs1`), and at least two redo log files (`redo1_1`, `redo1_2`).
 - For a RAC database, the file must specify one automatic undo tablespace datafile (`undotbsn`) and two redo log files (`redon_1`, `redon_2`) for each instance.
 - Specify at least two control files (`control1`, `control2`).
 - To use manual instead of automatic undo management, specify a single RBS tablespace datafile (`rbs`) instead of the automatic undo management tablespaces.
5. Save the file and note the file name that you specified.
 6. When you are configuring the `oracle` user's environment later in this chapter, set the `DBCA_RAW_CONFIG` environment variable to specify the full path to this file.

Verify that the Required Software is Running

Before installing Oracle Real Application Clusters, verify that the `hangcheck-timer` module is loaded and configured correctly. If you are using Oracle Cluster File System for Linux, you must also verify that it is configured correctly. The following sections describe how to complete these tasks.

Note: Apart from the `hangcheck-timer` module, the cluster software mentioned in the following sections is not required to install Oracle Real Application Clusters. However, if this software is installed, Oracle Cluster Ready Services (CRS) can integrate with it.

Check the Configuration of the hangcheck-timer Module

The `hangcheck-timer` module monitors the Linux kernel for extended operating system hangs that could affect the reliability of a RAC node and cause a database corruption. If a hang occurs, the module reboots the node in seconds.

You can use the `hangcheck_tick` and `hangcheck_margin` parameters to control the behavior of the module, as follows:

- The `hangcheck_tick` parameter defines how often, in seconds, the `hangcheck-timer` checks the node for hangs. The default value is 60 seconds.
- The `hangcheck_margin` parameter defines how long the timer waits, in seconds, for a response from the kernel. The default value is 180 seconds.

If the kernel fails to respond within the sum of the `hangcheck_tick` and `hangcheck_margin` parameter values, the `hangcheck-timer` module reboots the system. Using the default values, the node would be rebooted if the kernel fails to respond within 240 seconds.

To verify that the `hangcheck-timer` module is running on every node:

1. Enter the following command on each node to determine which kernel modules are loaded:

```
# /sbin/lsmmod
```

2. If the `hangcheck-timer` module is not listed for any node, enter a command similar to the following to start the module on that node:

```
# /sbin/insmod hangcheck-timer hangcheck_tick=30 hangcheck_margin=180
```

3. To ensure that the module is loaded every time the system reboots, verify that the local system startup file contains the command shown in the previous step, or add it if necessary:

- Red Hat:

On Red Hat Enterprise Linux systems, add the command to the `/etc/rc.d/rc.local` file.

- UnitedLinux

On UnitedLinux systems, add the command to the `/etc/init.d/boot.local` file.

Check the Configuration of OCFS

If you intend to use Oracle Cluster File System for database files, verify that it is configured correctly, as follows:

Note: For more information about configuring OCFS, see the documentation provided with the kit or available from the following Web site:

<http://oss.oracle.com/projects/ocfs>

1. Enter the following command to verify that OCFS is configured to start at runlevels 2, 3, and 5:

```
# /usr/sbin/chkconfig --list ocfs
```

2. Verify that the OCFS file systems you want to use are mounted and are specified in the `/etc/fstab` file.

Stop Existing Oracle Processes

Caution: If you are installing additional Oracle Database 10g products in an existing Oracle home, stop **all** processes running in the Oracle home. You must complete this task to enable the Installer to relink certain executables and libraries.

If you choose to create a database during the installation, most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the IPC key value EXTPROC. However, if an existing Oracle Net listener process is using the same port or key value, the Installer can only configure the new listener; it cannot start it. To ensure that the new listener process starts during the installation, you must shut down any existing listeners before starting the Installer.

To determine whether an existing listener process is running and to shut it down if necessary, follow these steps:

1. Switch user to `oracle`:

```
# su - oracle
```

2. Enter the following command to determine whether a listener process is running and to identify its name and the Oracle home directory in which it is installed:

```
$ ps -ef | grep tnslnsr
```

This command displays information about the Oracle Net listeners running on the system:

```
... oracle_home1/bin/tnslnsr LISTENER -inherit
```

In this example, `oracle_home1` is the Oracle home directory where the listener is installed and `LISTENER` is the listener name.

3. Set the `ORACLE_HOME` environment variable to specify the appropriate Oracle home directory for the listener:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=oracle_home1
$ export ORACLE_HOME
```

- C or tcsh shell:

```
$ setenv ORACLE_HOME oracle_home1
```

4. Enter the following command to identify the TCP/IP port number and IPC key value that the listener is using:

```
$ $ORACLE_HOME/bin/lsnrctl status listenername
```

Note: If the listener uses the default name `LISTENER`, you do not have to specify the listener name in this command.

5. Enter a command similar to the following to stop the listener process:

```
$ $ORACLE_HOME/bin/lsnrctl stop listenername
```

6. Repeat this procedure to stop all listeners running on this system and on all other cluster nodes.

Configure the oracle User's Environment

You run the Installer from the `oracle` account. However, before you start the Installer you must configure the environment of the `oracle` user. To configure the environment, you must:

- Set the default file mode creation mask (`umask`) to `022` in the shell startup file.
- Set the `DISPLAY` and `ORACLE_BASE` environment variables.

To set the `oracle` user's environment, follow these steps:

1. Start a new terminal session, for example, an X terminal (`xterm`).
2. Enter the following command to ensure that X Window applications can display on this system:

```
$ xhost +
```

3. If you are not already logged in to the system where you want to install the software, log in to that system as the `oracle` user.
4. If you are not logged in as the `oracle` user, switch user to `oracle`:

```
$ su - oracle
```

5. To determine the default shell for the `oracle` user, enter the following command:

```
$ echo $SHELL
```

6. Open the `oracle` user's shell startup file in any text editor:

Note: On Red Hat Linux, `.bash_profile` is the user startup file for the Bash shell.

- Bourne shell (`sh`), Bash shell (`bash`), or Korn shell (`ksh`):

```
$ vi .profile
```

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

```
% vi .login
```

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

```
umask 022
```

8. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, remove the appropriate lines from the file.

9. Save the file and exit from the editor.

10. To run shell startup script, enter one of the following command:

- Bash shell on Red Hat Linux:

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

- Bourne, Bash, or Korn shell:

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

- C shell:

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

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

- Bourne, Bash, or Korn shell:

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

- C shell:

```
$ setenv DISPLAY local_host:0.0
```

In this example, *local_host* is the host name or IP address of the system you want to use to display the Installer (your workstation or PC).

12. If you determined that the `/tmp` directory has less than 200 MB of free disk space, identify a file system with at least 200 MB of free space and set the `TEMP` and `TMPDIR` environment variables to specify a temporary directory on this file system:

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

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

```
$ su - root
# mkdir /mount_point/tmp
# chmod a+wr /mount_point/tmp
# exit
```

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

- * Bourne, Bash, or Korn shell:

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

- * C shell:

```
$ setenv TEMP /mount_point/tmp
$ setenv TMPDIR /mount_point/tmp
```

13. Enter commands similar to the following to set the `ORACLE_BASE` environment variable:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle
$ export ORACLE_BASE
```

- C shell:

```
% setenv ORACLE_BASE /u01/app/oracle
```

In these examples, `/u01/app/oracle` is the Oracle base directory that you created earlier.

14. If necessary, set the following environment variables. If you must specify several values for an environment variable, for example `PATH`, separate the values with a colon (`:`).

Environment Variable	Required By	Sample Setting and Description
<code>DBCA_RAW_CONFIG</code>	Oracle Database installations that use raw devices for database storage.	<code>\$ORACLE_BASE/oradata/dbname/dbname_raw.conf</code> Specifies the location of the raw device mapping file
<code>COBDIR</code>	Pro*COBOL	<code>/opt/lib/cobol</code> Specifies the directory where COBOL is installed on the system. Note: To install Pro*COBOL, you must choose a custom software installation.
<code>PATH</code>	Pro*COBOL	<code>\$COBDIR/bin:\$PATH</code> Specifies the directory where the COBOL compiler executable is located.

15. Enter the following command to ensure that the `ORACLE_HOME` and `TNS_ADMIN` environment variable are not set:

```
$ unset ORACLE_HOME
$ unset TNS_ADMIN
```

Note: If the `ORACLE_HOME` environment variable is set, the Installer uses the value it specifies as the default path for the Oracle home directory. However, if you set the `ORACLE_BASE` environment variable, Oracle recommends that you unset this environment variable and choose the default path suggested by the Installer.

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

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

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

Pre-Installation Tasks for Installing RAC on Solaris Operating System (SPARC 64-Bit)

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer. It includes information about the following tasks:

- Log In to the System as root
- Check the Hardware Requirements
- Check the Network Requirements
- Check the Software Requirements
- Create Required UNIX Groups and User
- Configure Kernel Parameters
- Identify Required Software Directories
- Identify or Create an Oracle Base Directory
- Configure Disk Storage for Oracle CRS and Database Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Configure Raw Logical Volumes
- Verify that the Cluster Software is Configured and Running
- Stop Existing Oracle Processes
- Configure the oracle User's Environment

Log In to the System as root

Before you install the Oracle software, you must complete several tasks as the `root` user. To log in as the `root` user, complete one of the following procedures:

Note: Unless you intend to complete a silent installation, you must install the software from an X Window System workstation, an X terminal, or a PC or other system with X server software installed.

For more information about non-interactive and silent installations, see the appendices in *Oracle Database Installation Guide for UNIX Systems*.

- If you are installing the software from an X Window System workstation or X terminal:
 1. Start a local terminal session, for example, an X terminal (`xterm`).
 2. If you are not installing the software on the local system, enter the following command to enable remote hosts to display X applications on the local X server:

```
$ xhost +
```

3. If you want to install the software on a remote system, enter a command similar to the following to connect to that system:

```
$ telnet remote_host
```

4. If you are not logged in as the `root` user, enter the following command to switch user to `root`:

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

- If you are installing the software from a PC or other system with X server software installed:

Note: If necessary, see your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

1. Start the X server software.
2. Configure the security settings of the X server software to permit remote hosts to display X applications on the local system.
3. Connect to the remote system where you want to install the software and start a terminal session on that system, for example, an X terminal (`xterm`).
4. If you are not logged in as the `root` user on the remote system, enter the following command to switch user to `root`:

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

Check the Hardware Requirements

The system must meet the following minimum hardware requirements:

- 512 MB of physical RAM
- 1 GB of swap space (or twice the size of RAM)
On systems with 2 GB or more of RAM, the swap space can be between one and two times the size of RAM.
- 400 MB of disk space in the `/tmp` directory
- Up to 4 GB of disk space for the Oracle software, depending on the installation type
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirements for databases that use Automatic Storage Management (ASM) or raw device storage is described later in this chapter.

Additional disk space, either on a file system or in an ASM disk group, is required for the flash recovery area if you choose to configure automated backups.

To ensure that the system meets these requirements, follow these steps:

1. To determine the physical RAM size, enter the following command:

```
# /usr/sbin/prtconf | grep "Memory size"
```

 If the size of the physical RAM installed in the system is less than the required size, you must install more memory before continuing.
2. To determine the size of the configured swap space, enter the following command:

```
# /usr/sbin/swap -s
```

 If necessary, see your operating system documentation for information about how to configure additional swap space.
3. To determine the amount of disk space available in the `/tmp` directory, enter the following command:

```
# df -k /tmp
```

If there is less than 400 MB of disk space available in the `/tmp` directory, complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to achieve the required disk space.
 - Set the `TEMP` and `TMPDIR` environment variables when setting the `oracle` user's environment (described later).
 - Extend the file system that contains the `/tmp` directory. If necessary, contact your system administrator for information about extending file systems.
4. To determine the amount of free disk space on the system, enter the following command:

```
# df -k
```

5. To determine whether the system architecture can run the software, enter the command:

```
# /bin/isainfo -kv
```

The expected output is:

```
64-bit sparcv9 kernel modules
```

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

Check the Network Requirements

Check that you have the networking hardware and internet protocol (IP) addresses required for an Oracle Real Application Clusters installation.

Note: For the most up-to-date information about supported network protocols and hardware for RAC installations, see the *Certify* pages on the Oracle *MetaLink* Web site:

<http://metalink.oracle.com>

Network Hardware Requirements

Each node in the cluster must meet the following requirements:

- Each node must have at least two network adapters; one for the public network interface and one for the private network interface (the interconnect).
- The interface names associated with the network adapters for each network must be the same on all nodes.
- For increased reliability, you can configure redundant public and private network adapters for each node.
- For the public network, each network adapter must support TCP/IP.
- For the private network, the interconnect must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended).

Note: UDP is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS.

Token-Ring is not supported for the interconnect.

- For the private network, Oracle supports the following interconnect protocols and hardware:
 - User datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended)

Note: UDP is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS.

- Hyper Messaging protocol (HMP) using HyperFabric adapters and switches
- For the private network, Oracle supports the following interconnect protocols and hardware:
 - User datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended)

Note: UDP is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS.

- Remote shared memory (RMS) using peripheral component interconnect-scalable coherent interface (PCI-SCI) adapters and SCI switches
- For the private network, Oracle supports the following interconnect protocols and hardware:
 - Reliable datagram (RDG) using Memory Channel adapters and hubs
 - RDG or user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended)

Note: RDG is the default interconnect protocol for RAC and TCP is the interconnect protocol for Oracle CRS. To use UDP as the interconnect protocol for RAC, you must relink the `oracle` executable. See the *Oracle Database Administrator's Reference for UNIX Systems* for more information.

IP Address Requirements

Before starting the installation, you must identify or obtain the following IP addresses for each node:

- An IP address and an associated host name registered in the domain name service (DNS) for each public network interface
- One unused virtual IP address and an associated virtual host name registered in DNS that you will configure for the primary public network interface

The virtual IP address must be in the same subnet as the associated public interface. After installation, you can configure clients to use the virtual host name or IP address. If a node fails, its virtual IP address fails over to another node.

- A private IP address and optional host name for each private interface

Oracle recommends that you use private network IP addresses for these interfaces, for example: 10.*.* or 192.168.*.*. You can use the `/etc/hosts` file on each node to associate private host names with private IP addresses.

For example, if each node has two public and two private interfaces, you might have the following host names and IP addresses on one of the nodes (`rac1`) and similar host names and IP addresses on the other nodes:

Host Name	Type	IP Address	Registered In
<code>rac1.mydomain.com</code>	Public	143.47.43.100	DNS
<code>rac1-2.mydomain.com</code>	Public	143.46.51.101	DNS
<code>rac1-vip.mydomain.com</code>	Virtual	143.46.43.104	DNS

Host Name	Type	IP Address	Registered In
rac1-priv1	Private	10.0.0.1	/etc/hosts
rac1-priv2	Private	10.0.0.2	/etc/hosts

Checking the Network Requirements

To verify that each node meets the requirements, follow these steps:

1. If necessary, install the network adapters for the public and private networks and configure them with either public or private IP addresses.
2. Register the host names and IP addresses for the public network interfaces in DNS.
3. For each node, register one virtual host name and IP address in DNS.
4. For each private interface on every node, add a line similar to the following to the `/etc/hosts` file on all nodes, specifying the private IP address and associated private host name:

```
10.0.0.1    rac1-priv1
```

5. To identify the interface name and associated IP address for every network adapter, enter the following command:

```
# /usr/sbin/ifconfig -a
```

From the output, identify the interface name and IP address for all network adapters that you want to specify as public or private network interfaces.

Note: When you install Oracle CRS and RAC, you will require this information.

Check the Software Requirements

See the following section for information about checking the software requirements.

Note: The Oracle Universal Installer performs checks on your system to verify that it meets the requirements listed for your platform. To ensure that these checks pass, verify the requirements before you start the Installer.

- ["Check for Required Software"](#) on page 6-6
- ["Check for Required Patches"](#) on page 6-8

Check for Required Software

Depending on the products that you intend to install, verify that the following software is installed on the system. The procedure following the table describes how to check these requirements.

Installation Type or Product	Requirement
All installations	<p>Operating system versions:</p> <p>Solaris 8 or Solaris 9, 64-bit</p> <hr/> <p>Operating system packages:</p> <p>SUNWarc SUNWbtool SUNWhea SUNWlibm SUNWlibms SUNWsprout SUNWsprrox SUNWtoo SUNWllof SUNWl1cs SUNWl15cs SUNWxwEnt</p> <p>Note: You might also require additional font packages for Java, depending on your locale. See the following Web site for more information:</p> <p>http://java.sun.com/j2se/1.4.2/font-requirements.html</p>
Oracle Spatial	<p>Motif RunTime Kit:</p> <p>SUNWmfrun</p> <p>Note: This software is required only to build the sample programs.</p>
Real Application Clusters	Sun Cluster 3.1
Oracle Messaging Gateways	<p>IBM WebSphere MQ V5.3, client and server:</p> <p>mqm</p> <hr/> <p>TIBCO Rendezvous 7.2</p>
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Sun ONE Studio 8 (C and C++ 5.5)
Oracle JDBC/OCI Drivers	<p>You can use the following optional JDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation:</p> <ul style="list-style-type: none"> ■ Sun Java 2 SDK Standard Edition 1.3.1_10 and the JNDI extension ■ Sun Java 2 SDK Standard Edition 1.2.2_17 and the JNDI extension <p>Note: JDK 1.4.2_01 is installed with this release.</p>
Oracle Transparent Gateway for Sybase	<p>Either of the following:</p> <ul style="list-style-type: none"> ■ Sybase Adaptive Server Enterprise V12 or 12.5 ■ Sybase Open Client V12 or 12.5 <p>Open Client is required if the Sybase server is not on the local node.</p>

Installation Type or Product	Requirement
Oracle Transparent Gateway for Informix	<p>One of the following:</p> <ul style="list-style-type: none"> ■ Informix Dynamic Server v7.3, 7.4, or 9.* ■ Informix ESQL/C v7.3, 7.4, or 9.* <p>ESQL/C is required if the Informix server is not on the local node.</p> <p>Note: This version of the gateway does not support Informix SE.</p>
Oracle Transparent Gateway for Ingress	Ingres II version 2.0, 2.5, or 2.6
Oracle Transparent Gateway for Teradata	NCR Teradata ODBC Driver v02.04.00.00

To ensure that the system meets these requirements, follow these steps:

1. To determine which version of Solaris is installed, enter the following command:

```
# uname -r
```

If the operating system version is lower than Solaris 8 (5.8), upgrade your operating system to this level.

2. To determine whether the required packages are installed, enter a command similar to the following:

```
# pkginfo -i SUNWarc SUNWbtool SUNWhea SUNWlibm SUNWlibms SUNWsprot \
SUNWsprox SUNWtoo SUNWilof SUNWilcs SUNWil5cs SUNWxwfont
```

If a package is not installed, then install it. See your operating system or software documentation for information about installing packages.

Check for Required Patches

Depending on the products that you intend to install, verify that the following patches are installed on the system. The procedure following the table describes how to check these requirements.

Note: The patch versions shown in the following table are minimum versions. Higher versions of the same patches are also supported.

Installation Type or Product	Requirement
All installations	<p>Patches for Solaris 8:</p> <p>All of the patches included in the J2SE Patch Cluster for Solaris 8:</p> <ul style="list-style-type: none"> ■ 108528-23, SunOS 5.8: kernel update patch ■ 108652-66, X11 6.4.1: Xsun patch ■ 108773-18, SunOS 5.8: IIM and X I/O Method patch ■ 108921-16, CDE 1.4: dtwm patch ■ 108940-53, Motif 1.2.7 and 2.1.1: Runtime lib. patch for Solaris 8 ■ 108987-13, SunOS 5.8: Patch for patchadd and patchrm ■ 108989-02, /usr/kernel/sys/acctctl & /.../exacctsyst patch ■ 108993-19, SunOS 5.8: LDAP2 client, libc, ... lib. patch ■ 109147-24, SunOS 5.8: linker patch ■ 110386-03, SunOS 5.8: RBAC Feature Patch ■ 111023-02, SunOS 5.8: /kernel/fs/mntfs and ... sparcv9/mntfs ■ 111111-03, SunOS 5.8: /usr/bin/nawk patch ■ 111308-03, SunOS 5.8: /usr/lib/libmtmalloc.so.1 patch ■ 111310-01, SunOS 5.8: /usr/lib/libdhcagent.so.1 patch ■ 112396-02, SunOS 5.8: /usr/bin/fgrep patch <p>The following additional patches:</p> <ul style="list-style-type: none"> ■ 111721-04, SunOS 5.8: Math Library (libm) patch ■ 112003-03, SunOS 5.8: Unable to load fontset ... iso-1 or iso-15 ■ 112138-01, SunOS 5.8: usr/bin/domainname patch
All installations	<p>Patches for Solaris 9:</p> <ul style="list-style-type: none"> ■ 112233-11: SunOS 5.9: Kernel Patch ■ 111722-04: SunOS 5.9: Math Library (libm) patch
Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Patches for Solaris 8 and Solaris 9:</p> <ul style="list-style-type: none"> ■ 112758-02, dbx 7.1: Patch for S1S8CC Debugger ■ 112760-05, C 5.5: Patch for S1S8CC C compiler ■ 112763-06, Compiler Common S1S8CC: Patch for S1S8CC C C++ F77 F95 ■ 113817-03, C++ 5.5: Patch for S1S8CC C++ compiler ■ 113820-02, Compiler Common 7.1: Patch for S1S8CC Math Libraries ■ 113823-03, ILD 4.1: Patch for S1S8CC Incremental Linker
Real Application Clusters	<p>Sun Cluster patches for Solaris 8:</p> <ul style="list-style-type: none"> ■ 111488-08, Sun Cluster 3.0/3.1: mediator patch ■ 113800-06, Sun Cluster 3.1: Core/Sys Admin Patch <p>To obtain the following patches, contact your local Sun solution center:</p> <p>114492-02, 115058-03, 115062-02, 115064-01, 115066-01, 115070-01, 115072-01, 115076-01, 115078-01, 115080-01, 115082-01, 115084-01, 115086-01, 115089-01</p> <p>Sun Cluster patches for Solaris 9:</p> <p>113801: Sun Cluster 3.1: Core/Sys Admin Patch</p>

Installation Type or Product	Requirement
Oracle Messaging Gateway	Corrective service diskettes (CSDs) for WebSphere MQ: CSD06 or later for WebSphere MQ V5.3

To ensure that the system meets these requirements, follow these steps:

1. To determine whether an operating system patch is installed, enter a command similar to the following:

```
# /usr/sbin/patchadd -p | grep patch_number
```

If an operating system patch is not installed, download it from the following Web site and install it:

<http://sunsolve.sun.com>

2. If you require a CSD for WebSphere MQ, see the following Web site for download and installation information:

<http://www.ibm.com/software/integration/mqfamily/support/summary/sun.html>

Create Required UNIX Groups and User

Depending on whether this is the first time Oracle software is being installed on this system and on the products that you are installing, you may need to create several UNIX groups and a UNIX user.

The following UNIX groups and user are required if you are installing Oracle Database:

- The OSDBA group (`dba`)

You must create this group the first time you install Oracle Database software on the system. It identifies UNIX users that have database administrative privileges (the SYSDBA privilege). The default name for this group is `dba`.

If you want to specify a group name other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of this group. In this case, the Installer prompts you to specify the name of this group.

- The OSOPER group (`oper`)

This is an optional group. Create this group if you want a separate group of UNIX users to have a limited set of database administrative privileges (the SYSOPER privilege). By default, members of the OSDBA group also have the SYSOPER privilege.

If you want to specify a separate OSOPER group, other than the default `dba` group, you must choose the Custom installation type to install the software or start the Installer as a user that is not a member of the `dba` group. In this case, the Installer prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

- An unprivileged user (`nobody`)

You must verify that the unprivileged user `nobody` exists on the system. The `nobody` user must own the external jobs (`ext job`) executable after the installation.

The following UNIX group and user are required for all installations:

- The Oracle Inventory group (`oinstall`)

You must create this group the first time you install Oracle software on the system. The usual name chosen for this group is `oinstall`. This group owns the Oracle inventory, which is a catalog of all Oracle software installed on the system.

Note: If Oracle software is already installed on the system, the existing Oracle Inventory group must be the primary group of the UNIX user that you use to install new Oracle software. The following sections describe how to identify an existing Oracle Inventory group.

- The Oracle software owner user (`oracle`)

You must create this user the first time you install Oracle software on the system. This user owns all of the software installed during the installation. The usual name chosen for this user is `oracle`. This user must have the Oracle Inventory group as its primary group. It must also have the OSDBA and OSOPER groups as a secondary groups.

A single Oracle Inventory group is required for all installations of Oracle software on the system. After the first installation of Oracle software, you must use the same Oracle Inventory group for all subsequent Oracle software installations on that system. However, you can choose to create different Oracle software owner users, OSDBA groups, and OSOPER groups (other than `oracle`, `dba`, and `oper`) for separate installations. By using different groups for different installations, members of these different groups have DBA privileges only on the associated databases rather than on all databases on the system.

See Also: *The Oracle Database Administrator's Reference for UNIX Systems* and the *Oracle Database Administrator's Guide* contain more information about the OSDBA and OSOPER groups and the SYSDBA and SYSOPER privileges.

The following sections describe how to create the required UNIX user and groups.

Note: The following sections describe how to create local users and groups. As an alternative to creating local users and groups, you could create the appropriate users and groups in a directory service, for example, Network Information Services (NIS). For information about using directory services, contact your system administrator or see your operating system documentation.

Creating the Oracle Inventory Group

You must create the Oracle Inventory group if it does not already exist. The following subsections describe how to determine the Oracle Inventory group name, if it exists, and how to create it if necessary.

Determine Whether the Oracle Inventory Group Exists

When you install Oracle software on the system for the first time, the Installer creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group and the path of the Oracle Inventory directory. To determine whether the Oracle Inventory group exists, enter the following command:

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

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inst_group` parameter shows the name of the Oracle Inventory group (`oinstall`).

Create the Oracle Inventory Group

If the `oraInst.loc` file does not exist, create the Oracle Inventory group, `oinstall`, by entering the following command:

```
# /usr/sbin/groupadd oinstall
```

Creating the OSDBA Group

You must create an OSDBA group in the following circumstances:

- An OSDBA group does not exist, for example, if this is the first installation of Oracle Database software on the system
- An OSBDA group exists, but you want to give a different group of UNIX users database administrative privileges in a new Oracle installation

If the OSDBA group does not exist or if you require a new OSDBA group, create it as follows, using the group name `dba` unless a group with that name already exists:

```
# /usr/sbin/groupadd dba
```

Creating an OSOPER Group (Optional)

Create an OSOPER group only if you want to identify a group of UNIX users with a limited set of database administrative privileges (SYSOPER operator privileges). For most installations, it is sufficient to create only the OSDBA group. If you want to use an OSOPER group, you must create it in the following circumstances:

- If an OSOPER group does not exist, for example, if this is the first installation of Oracle Database software on the system
- If an OSOPER group exists, but you want to give a different group of UNIX users database operator privileges in a new Oracle installation

If you require a new OSOPER group, create it as follows, using the group name `oper` unless a group with that name already exists:

```
# /usr/sbin/groupadd oper
```

Creating the Oracle Software Owner User

You must create an Oracle software owner user in the following circumstances:

- If an Oracle software owner user does not exist, for example, if this is the first installation of Oracle software on the system
- If an Oracle software owner user exists, but you want to use a different UNIX user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation

Determine Whether an Existing Oracle Software Owner User Exists

To determine whether an Oracle software owner user named `oracle` exists, enter the following command:

```
# id -a oracle
```

If the `oracle` user exists, the output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

If the user exists, determine whether you want to use the existing user or create a new user. If you want to use the existing user, ensure that the user's primary group is the Oracle Inventory group and that it is a member of the appropriate OSDBA and OSOPER groups. See one of the following sections for more information:

Note: If necessary, contact your system administrator before using or modifying an existing user.

- If you want to use the existing Oracle software owner user, and the user's primary group is the Oracle Inventory group, see the ["Verifying That the UNIX User nobody Exists"](#) on page 6-14
- To modify an existing user, see the ["Modify an Existing Oracle Software Owner User"](#) section on page 6-14.
- To create a new user, see the following section.

Create a New Oracle Software Owner User

If the Oracle software owner user does not exist or if you require a new Oracle software owner user, create it as follows, using the user name `oracle` unless a user with that name already exists.

1. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example, `oinstall`
- The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba , oper`

2. Set the password of the `oracle` user:

```
# passwd oracle
```

See the section to continue. ["Verifying That the UNIX User nobody Exists"](#) on page 6-14

Modify an Existing Oracle Software Owner User

If the `oracle` user exists, but its primary group is not `oinstall` or it is not a member of the appropriate `OSDBA` or `OSOPER` groups, you can modify it by entering a command similar to the following, specifying the primary group using the `-g` option and any required secondary groups using the `-G` option:

```
# /usr/sbin/usermod -g oinstall -G dba[,oper] oracle
```

Verifying That the UNIX User `nobody` Exists

Before installing the software, verify that the UNIX user `nobody` exists on the system:

1. To determine whether the user exists, enter the following command:

```
# id nobody
```

If this command displays information about the `nobody` user, you do not have to create that user.

2. If the `nobody` user does not exist, enter the following command to create it:

```
# /usr/sbin/useradd nobody
```

3. Repeat this procedure on all of the other cluster nodes.

Create Identical Users and Groups on Other Cluster Nodes

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

The Oracle software owner user and the Oracle Inventory, `OSDBA`, and `OSOPER` groups must exist and be identical on all cluster nodes. To create these identical users and groups, you must identify the user ID and group IDs assigned them on the node where you created them, then create the user and groups with the same name and ID on the other cluster nodes.

Identifying the User and Group IDs

To determine the user ID (UID) of the Oracle software owner user and the group IDs (GID) of the Oracle Inventory, `OSDBA`, and `OSOPER` groups, follow these steps:

1. Enter the following command:

```
# id -a oracle
```

The output from this command is similar to the following:

```
uid=440(oracle) gid=200(oinstall) groups=201(dba),202(oper)
```

2. From the output, identify the UID for the `oracle` user and the GIDs for the groups to which it belongs.

Create the User and Groups on the Other Cluster Nodes

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

1. Log in to the next cluster node as `root`.
2. Enter commands similar to the following to create the `oinstall` and `dba` groups, and if required, the `oper` group. Use the `-g` option to specify the correct GID for each group.

```
# /usr/sbin/groupadd -g 200 oinstall
# /usr/sbin/groupadd -g 201 dba
# /usr/sbin/groupadd -g 202 oper
```

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

3. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -u 200 -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-u` option specifies the user ID, which must be the user ID that you identified in the previous subsection.
- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`
- The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba , oper`

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

4. Set the password of the `oracle` user:

```
# passwd oracle
```

Configure SSH on All Cluster Nodes

Before you install and use Oracle Real Application clusters, you must configure secure shell (SSH) for the `oracle` user on all cluster nodes. The Installer uses the `ssh` and `scp` commands during installation to run remote commands on and copy files to the other cluster nodes. You must configure SSH so that these commands do not prompt for a password.

Note: This section describes how to configure OpenSSH version 3. If SSH is not available, the Installer attempts to use `rsh` and `rcp` instead.

To configure SSH, complete the following steps on each cluster node:

1. Log in as the `oracle` user.
2. If necessary, create the `.ssh` directory in the `oracle` user's home directory and set the correct permissions on it:

```
$ mkdir ~/.ssh
$ chmod 755 ~/.ssh
```

3. Enter the following commands to generate an RSA key for version 2 of the SSH protocol:

```
$ /usr/bin/ssh-keygen -t rsa
```

At the prompts:

- Accept the default location for the key file
- Enter and confirm a pass phrase that is different from the `oracle` user's password

This command writes the public key to the `~/.ssh/id_rsa.pub` file and the private key to the `~/.ssh/id_rsa` file. Never distribute the private key to anyone.

4. Enter the following commands to generate a DSA key for version 2 of the SSH protocol:

```
$ /usr/bin/ssh-keygen -t dsa
```

At the prompts:

- Accept the default location for the key file
- Enter and confirm a pass phrase that is different from the `oracle` user's password

This command writes the public key to the `~/.ssh/id_dsa.pub` file and the private key to the `~/.ssh/id_dsa` file. Never distribute the private key to anyone.

5. Copy the *contents* of the `~/.ssh/id_rsa.pub` and `~/.ssh/id_dsa.pub` files to the `~/.ssh/authorized_keys` file on this node and to the same file on all other cluster nodes.

Note: The `~/.ssh/authorized_keys` file on every node must contain the contents from all of the `~/.ssh/id_rsa.pub` and `~/.ssh/id_dsa.pub` files that you generated on all cluster nodes.

6. Change the permissions on the `~/.ssh/authorized_keys` file on all cluster nodes:

```
$ chmod 644 ~/.ssh/authorized_keys
```

At this point, if you use `ssh` to log in to or run a command on another node, you are prompted for the pass phrase that you specified when you created the DSA key.

To enable the Installer to use the `ssh` and `scp` commands without being prompted for a pass phrase, follow these steps:

1. On the system where you want to install the software, log in as the `oracle` user.
2. Enter the following commands:

```
$ exec /usr/bin/ssh-agent $SHELL
$ /usr/bin/ssh-add
```

3. At the prompts, enter the pass phrase for each key that you generated.

If you have configured SSH correctly, you can now use the `ssh` or `scp` commands without being prompted for a password or a pass phrase.

4. To test the SSH configuration, enter the following commands from the same terminal session, testing the configuration of each cluster node:

```
$ ssh nodename1 date
$ ssh nodename2 date
.
.
```

These commands should display the date set on each node. If any node prompts for a password or pass phrase, verify that the `~/.ssh/authorized_keys` file on that node contains the correct public keys.

Note: The first time you use SSH to connect to a node from a particular system, you might see a message stating that the authenticity of the host could not be established. Enter `yes` at the prompt to continue. You should not see this message again when you connect from this system to that node.

5. You must run the Installer from this session or remember to repeat steps 2 and 3 before you start the Installer from a different terminal session.

Set Up User Equivalence for `rsh` and `rcp` on All Cluster Nodes

Note: This section describes how to set up user equivalence for `rcp`, which the Installer uses when copying Oracle software to the other cluster nodes. If you prefer, you can configure the Secure Shell (SSH) tool suite, so that the Installer uses `scp` instead of `rcp`. See the SSH documentation for information about setting up user equivalence for `scp`.

Before you install and use Oracle Real Application clusters, you must set up user equivalence on all cluster nodes, as follows:

1. On all cluster nodes, edit the `/etc/hosts.equiv` file in any text editor and add entries similar to the following for all nodes in the cluster:

```
node1-public oracle
node1-public.domain oracle
node2-public oracle
node2-public.domain oracle
.
.
.
```

In this example:

- `noden_public` is the host name associated with the primary public network interface
- `domain` is the domain name of the system, for example `mydomain.com`

Note: Alternatively, you can create a `.rhosts` file using the same format in the `oracle` user's home directory on all nodes in the cluster.

- To test user equivalence on all cluster nodes, including the first node, enter commands similar to the following from the node where you will run the Installer:

```
# su oracle
$ rsh nodename1 date
$ rsh nodename2 date
.
.
.
$ exit
```

Each command should return only the date setting on the specified node. If a node prompts for a password or displays any other text, then user equivalence is not set up correctly on that node. In this case, verify that the `/etc/hosts.equiv` (or `~oracle/.rhosts`) file has the correct entries.

Configure Kernel Parameters

Notes: The kernel parameter and shell limit values shown in the following sections are recommended values only. For production database systems, Oracle recommends that you tune these values to optimize the performance of the system. See your operating system documentation for more information about tuning kernel parameters.

You must set the kernel parameters and shell limits on all cluster nodes.

Verify that the kernel parameters shown in the following table are set to values greater than or equal to the recommended value shown. The procedure following the table describes how to verify and set the values.

Parameter	Recommended Value
<code>noexec_user_stack</code> ¹	1
<code>semsys:semaphore_semni</code>	100
<code>semsys:semaphore_semmns</code>	1024
<code>semsys:semaphore_semmsl</code>	256
<code>semsys:semaphore_semvmx</code>	32767
<code>shmsys:shmshmmax</code>	4294967295
<code>shmsys:shmshmmin</code> ¹	1
<code>shmsys:shmshmni</code>	100
<code>shmsys:shmshmseg</code> ¹	10

¹ This parameter is obsolete in Solaris 9

Note: If the current value for any parameter is higher than the value listed in this table, do not change the value of that parameter.

To view the current value specified for these kernel parameters, and to change them if necessary, follow these steps:

1. To view the current values of these parameters, enter the following commands:

```
# grep noexec_user_stac /etc/system
# /usr/sbin/sysdef | grep SEM
# /usr/sbin/sysdef | grep SHM
```

2. If you must change any of the current values, follow these steps:

- a. Create a backup copy of the `/etc/system` file, for example:

```
# cp /etc/system /etc/system.orig
```

- b. Open the `/etc/system` file in any text editor and, if necessary, add lines similar to the following (edit the lines if the file already contains them):

```
set noexec_user_stack=1
set semsys:seminfo_semmni=100
set semsys:seminfo_semmns=1024
set semsys:seminfo_semmsl=256
set semsys:seminfo_semvmx=32767
set shmsys:shminfo_shmmax=4294967295
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=100
set shmsys:shminfo_shmseg=10
```

- c. Enter the following command to reboot the system:

```
# /usr/sbin/reboot
```

- d. When the system restarts, log in and switch user to `root`.

3. Repeat this procedure on all other cluster nodes.

Identify Required Software Directories

You must identify or create four directories for the Oracle software, as follows:

- Oracle base directory
- Oracle Inventory directory
- CRS home directory
- Oracle home directory

The following subsections describe the requirements for these directories.

Oracle Base Directory

The Oracle base directory acts as a top-level directory for Oracle software installations. It is analogous to the `C:\Oracle` directory used for Oracle software installations on Windows systems. On UNIX systems, the Optimal Flexible Architecture (OFA) guidelines recommend that you use a path similar to the following for the Oracle base directory:

```
/mount_point/app/oracle_sw_owner
```

- *mount_point* is the mount point directory for the file system that will contain the Oracle software.

The examples in this guide use `/u01` for the mount point directory. However, you could choose another mount point directory, `/oracle` or `/opt/oracle` for example.

- *oracle_sw_owner* is the UNIX user name of the Oracle software owner, for example `oracle`.

You can use the same Oracle base directory for more than one installation or you can create separate Oracle base directories for different installations. If different UNIX users install Oracle software on the same system, each user must create a separate Oracle base directory. The following example Oracle base directories could all exist on the same system:

```
/u01/app/oracle
/u01/app/orauser
/opt/oracle/app/oracle
```

The following sections describe how to identify existing Oracle base directories that might be suitable for your installation and how to create a new Oracle base directory if necessary.

Regardless of whether you create a new Oracle base directory or decide to use an existing one, you must set the `ORACLE_BASE` environment variable to specify the full path to this directory.

Note: The Oracle base directory must be on a local file system.

The Oracle base directory can be on a local file system or on an NFS file system on a certified NAS device. Do not create the Oracle base directory on an OCFS version 1 file system.

The Oracle base directory can be on a local file system, a supported cluster file system, or on an NFS file system on a certified NAS device.

The Oracle base directory can be on a local file system, or on a supported cluster file system.

Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system. It is required by, and shared by, all Oracle software installations on a single system. The first time you install Oracle software on a system, the Installer prompts you to specify the path to this directory. Oracle recommends that you choose the following path:

```
oracle_base/oraInventory
```

If the Oracle base directory is on a cluster file system, you must specify a path for the Oracle Inventory directory on a local file system, to enable all of the nodes to have separate inventories.

The Installer creates the directory that you specify and sets the correct owner, group, and permissions on it. You do not need to create it.

Note: All Oracle software installations rely on this directory. Make sure that you back it up regularly.

Do not delete this directory unless you have completely removed all Oracle software from the system.

CRS Home Directory

The CRS home directory is the directory where you choose to install the software for Oracle Cluster Ready Services. You must install CRS in a separate home directory. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
oracle_base/product/10.1.0/crs
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Oracle Home Directory

The Oracle home directory is the directory where you choose to install the software for a particular Oracle product. You must install different Oracle products, or different releases of the same Oracle product, in separate Oracle home directories. When you run the Installer, it prompts you to specify the path to this directory, as well as a name that identifies it. The directory that you specify must be a subdirectory of the Oracle base directory. Oracle recommends that you specify a path similar to the following for the Oracle home directory:

```
oracle_base/product/10.1.0/db_1
```

The Installer creates the directory path that you specify under the Oracle base directory. It also sets the correct owner, group, and permissions on it. You do not need to create this directory.

Identify or Create an Oracle Base Directory

Before starting the installation, you must either identify an existing Oracle base directory or if required, create a new one. This section contains information about the following:

- ["Identifying an Existing Oracle Base Directory"](#) on page 6-21
- ["Creating a New Oracle Base Directory"](#) on page 6-23

Note: You can choose to create a new Oracle base directory, even if other Oracle base directories exist on the system.

Identifying an Existing Oracle Base Directory

Existing Oracle base directories might not have paths that comply with OFA guidelines. However, if you identify an existing Oracle Inventory directory or existing Oracle home directories, you can usually identify the Oracle base directories, as follows:

- Identify an existing Oracle Inventory directory

Enter the following command on all cluster nodes to view the contents of the `oraInst.loc` file:

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

If the `oraInst.loc` file exists, the output from this command is similar to the following:

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

The `inventory_loc` parameter identifies the Oracle Inventory directory (`oraInventory`) on that system. The parent directory of the `oraInventory` directory is typically an Oracle base directory. In the previous example, `/u01/app/oracle` is an Oracle base directory.

- Identify existing Oracle home directories

Enter the following command on all cluster nodes to view the contents of the `oratab` file:

```
# more /var/opt/oracle/oratab
```

If the `oratab` file exists, it contains lines similar to the following:

```
*:/u03/app/oracle/product/10.1.0/db_1:N
*/opt/orauser/infra_904:N
*/oracle/9.2.0:N
```

The directory paths specified on each line identify Oracle home directories. Directory paths that end with the user name of the Oracle software owner that you want to use are valid choices for an Oracle base directory. If you intend to use the `oracle` user to install the software, you could choose one of the following directories from the previous example:

```
/u03/app/oracle
/oracle
```

Note: If possible, choose a directory path similar to the first (`/u03/app/oracle`). This path complies with the OFA guidelines.

Before deciding to use an existing Oracle base directory for this installation, make sure that it satisfies the following conditions:

- It should not be on the same file system as the operating system.
- It must have an identical path on all cluster nodes or it must be on a supported cluster file system.
If you are not using a cluster file system, create identical Oracle base directories on the other nodes.
- It must have an identical path on all cluster nodes or it must be on a supported cluster file system.
If you are not using a cluster file system, create identical Oracle base directories on the other nodes.
- It must have up to 4 GB of free disk space.
- It must have sufficient free disk space on all cluster nodes, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB, depending on the platform

Use the following command to determine the free disk space on the file system where the Oracle base directory is located:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the ["Creating a New Oracle Base Directory"](#) section on page 6-23.
- If an Oracle base directory exists and you want to use it, see the ["Configure Disk Storage for Oracle CRS and Database Files"](#) section on page 6-24.

When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify the directory you chose.

- If an Oracle base directory does not exist on the system or if you want to create a new Oracle base directory, see the following section.

Creating a New Oracle Base Directory

Before you create a new Oracle base directory, you must identify an appropriate file system with at least 4 GB of free disk space.

Before you create a new Oracle base directory, you must identify an appropriate file system with sufficient free disk space, as follows:

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 5 GB, depending on the platform

To identify an appropriate file system, follow these steps:

1. Use the `df -k` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has appropriate free space.

Note: The file system can be a local file system, a supported cluster file system, or an NFS file system on a certified NAS device.

The path to the Oracle base directory must be the same on all nodes.

3. Note the name of the mount point directory for the file system that you identified.

To create the Oracle base directory and specify the correct owner, group, and permissions for it, follow these steps:

1. Enter commands similar to the following to create the recommended subdirectories in the mount point directory that you identified and set the appropriate owner, group, and permissions on them:

```
# mkdir -p /mount_point/app/oracle_sw_owner
# chown -R oracle:oinstall /mount_point/app
# chmod -R 775 /mount_point/app
```

If the mount point you identified is `/u01` and `oracle` is the user name of the Oracle software owner, the recommended Oracle base directory path is as follows:

```
/u01/app/oracle
```

2. If necessary, repeat the commands listed in the previous step to create the same directory on the other cluster nodes.
3. When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify this directory.

Configure Disk Storage for Oracle CRS and Database Files

This section describes the storage options available to you for storing Oracle Cluster Ready Services files, Oracle Database Files, and optionally, Oracle Database recovery files. After you choose the storage method that you want to use for each file type, see the following sections for information about configuring the required storage.

Note: You do not have to use the same storage option for each type of file.

Before you install Oracle Cluster Ready Services, you must choose the storage option that you want to use for the Oracle Cluster Registry (100 MB) and CRS voting disk (20 MB). You cannot use Automatic Storage Management to store these files, because they must be accessible before any Oracle instance starts.

Apart from ASM, you can use any of the storage options listed in the next section for the Oracle CRS files.

Choosing a Storage Option for Oracle CRS Files

If you want to create a database during the installation, you must choose one of the following storage options for the database files:

- Supported cluster file system
- Automatic Storage Management
- NFS file system with Fujitsu PRIMECLUSTER and a certified NAS device
- Shared logical volumes
- Raw partitions

Note: For the most up-to-date information about supported storage options for RAC installations, see the Certify pages on the Oracle*MetaLink* Web site:

<http://metalink.oracle.com>

Special Considerations for Sun Cluster: Using VERITAS Volume Manager

If you are not using Automated Storage Management (ASM), then you must use VERITAS Volume Manager to create your volumes and then change the permissions and ownerships on the volumes to the correct user such as `oracle` user. Also make sure that all of the volumes on your disks have the same names and the same permissions on each node. More information about VERITAS Volume Manager is available at:

<http://www.veritas.com>

Make sure that the parent directory of the raw device is owned by `root` and that this directory does not have write permission for any user other than `root`.

Sun Cluster Private Host Names

When installing CRS, the Cluster Configuration screen asks you to enter the private name to be used to interconnect the nodes within your cluster. When CRS is installed on Sun Cluster, the private names that you specify on this screen must be the private host names assigned by Sun Cluster.

Sun Cluster assigns default private hostnames during the initial cluster installation. The default private host name has the format of `clusternode $nodeid$ -priv`. For example, for node ID 1, the default private host name is `clusternode1-priv`. For node ID 2, the default private host name is `clusternode2-priv`.

Do not store private host names in the host's database. An `nsswitch` utility, see `nsswitch.conf(4)`, performs all host name look-ups for private host names. Private hostnames can be displayed using the following command:

```
$ /usr/cluster/bin/scconf -p
```

Choosing a Storage Option for Oracle Database Recovery Files

Additionally, if you want to enable automated backups during the installation, you must choose one of the following storage options for recovery files (the flash recovery area):

- Supported cluster file system
- Automatic Storage Management
- NFS file system with Fujitsu PRIMECLUSTER and a certified NAS device

The storage option that you choose for recovery files can be the same as or different to the option you choose for the database files.

Configuring Disk Storage

For information about how to configure disk storage before you start the installation, see one of the following sections depending on your choice:

- To use a file system for Oracle CRS, database, or recovery file storage, see the ["Create Directories for Oracle CRS, Database, or Recovery Files"](#) section on page 6-26.
- To use ASM for database or recovery file storage, see the ["Configure Disks for Automatic Storage Management"](#) section on page 6-29.
- To use raw devices (partitions or logical volumes) for Oracle CRS or database file storage, see the ["Configure Raw Logical Volumes"](#) section on page 6-34.

Create Directories for Oracle CRS, Database, or Recovery Files

If you decide to place the Oracle CRS, database, or recovery files on a file system, use the following guidelines when deciding where to place them.

Guidelines for Placing Oracle CRS Files on a File System

The Installer does not suggest a default location for the Oracle Cluster Registry (OCR) or the Oracle CRS voting disk. If you choose to create these files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk.
- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.
- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.

Note: NAS storage is currently supported only if you are using Fujitsu PRIMECLUSTER and a certified NAS device.

- It must have at least 100 MB of free disk space for the OCR and 20 MB of free disk space for the CRS voting disk.
- For improved reliability, you should choose a file system on a highly available storage device, for example, a RAID device that implements mirroring.
- If you are placing the Oracle Cluster Ready Services software on a shared file system, you can use the same file system for these files.
- The `oracle` user must have write permissions to create the files in the path you specify.

Note: If you are upgrading from Oracle9i Release 2, you can continue to use the raw device or shared file that you used for the SRVM configuration repository instead of creating a new file for the OCR.

Guidelines for Placing Oracle Database Files on a File System

If you choose to place the Oracle database files on a file system, use the following guidelines when deciding where to place them:

- You must choose a shared file system, for example, a cluster file system on a shared disk.

- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.
- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.

Note: NAS storage is currently supported only if you are using Fujitsu PRIMECLUSTER and a certified NAS device.

- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.

This default location is not recommended for production databases.

- You can choose either a single file system or more than one file system to store the database files:
 - If you want to use a single file system, choose a file system on a physical device that is dedicated to the database.

For best performance and reliability, choose a RAID device or a logical volume on more than one physical device and implement the stripe-and-mirror-everything (SAME) methodology.
 - If you want to use more than one file system, choose file systems on separate physical devices that are dedicated to the database.

This method enables you to distribute physical I/O and create separate control files on different devices for increased reliability. It also enables you to fully implement the OFA guidelines. You must choose either the Advanced database creation option or the Custom installation type during the installation to implement this method.
- If you intend to create a preconfigured database during the installation, the file system (or file systems) that you choose must have at least 1.2 GB of free disk space.

For production databases, you must estimate the disk space requirement depending on the use you want to make of the database.
- For optimum performance, the file systems that you choose should be on physical devices that are used only by the database.
- The `oracle` user must have write permissions to create the files in the path you specify.

Guidelines for Placing Oracle Recovery Files on a File System

Note: You must choose a location for recovery files only if you intend to enable automated backups during the installation.

If you choose to place the Oracle recovery files on a file system, use the following guidelines when deciding where to place them:

- To prevent disk failure from making both the database files and the recovery files unavailable, place the recovery files in a file system on a different physical disk from the database files.

Note: Alternatively use an ASM disk group with a normal or high redundancy level for either or both file types.

- You must choose a shared file system, for example, a cluster file system on a shared disk.
- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.
- You must choose a shared file system, for example, a cluster file system on a shared disk or an NFS file system on a certified NAS device.

Note: NAS storage is currently supported only if you are using Fujitsu PRIMECLUSTER and a certified NAS device.

- The file system that you choose should have at least 2 GB of free disk space.
The disk space requirement is the default disk quota configured for the flash recovery area (specified by the DB_RECOVERY_FILE_DEST_SIZE initialization parameter).
If you choose the Custom installation type or the Advanced database configuration option, you can specify a different disk quota value. After you create the database, you can also use Oracle Enterprise Manager Grid Control or Database Control to specify a different value.
For more information about sizing the flash recovery area, see the *Oracle Backup and Recovery Basics* manual.
- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory. You can choose this path only if you are using an Oracle base directory that is on a shared file system.
This default location is not recommended for production databases.
- The `oracle` user must have write permissions to create the files in the path you specify.

Creating Required Directories

Note: You must complete this procedure only if you want to place the Oracle CRS, database, or recovery files on a separate file system to the Oracle base directory.

To create directories for the Oracle CRS, database, or recovery files on separate file systems to the Oracle base directory, follow these steps:

1. If necessary, configure the shared file systems that you want to use and mount them on each node.

Note: Make sure that the file systems are configured to mount automatically when a node reboots.

2. Use the `df -k` command to determine the free disk space on each mounted file system.
3. From the display, identify the file systems that you want to use:

File Type	File System Requirements
CRS files	Choose a file system with at least 120 MB of free disk space
Datafiles	Choose either: <ul style="list-style-type: none"> ■ A single file system with at least 1.2 GB of free disk space ■ Two or more file systems with at least 1.2 GB of free disk space in total
Recovery files	Choose a file system with at least 2 GB of free disk space.

If you are using the same file system for more than one type of file, add the disk space requirements for each type to determine the total disk space requirement.

4. Note the names of the mount point directories for the file systems that you identified.
5. Enter commands similar to the following to create the recommended subdirectories in each of the mount point directories and set the appropriate owner, group, and permissions on them:
 - CRS file directory:

```
# mkdir /mount_point/oracrs
# chown oracle:oinstall /mount_point/oracrs
# chmod 775 /mount_point/oracrs
```
 - Database file directory:

```
# mkdir /mount_point/oradata
# chown oracle:oinstall /mount_point/oradata
# chmod 775 /mount_point/oradata
```
 - Recovery file directory (flash recovery area):

```
# mkdir /mount_point/flash_recovery_area
# chown oracle:oinstall /mount_point/flash_recovery_area
# chmod 775 /mount_point/flash_recovery_area
```
6. If you also want to use ASM or raw devices for storage, see one of the following sections:
 - ["Configure Disks for Automatic Storage Management"](#) on page 6-29
 - ["Configure Raw Logical Volumes"](#) on page 6-34

Otherwise see the ["Verify that the Cluster Software is Configured and Running"](#) section on page 6-40.

Configure Disks for Automatic Storage Management

This section describes how to configure disks for use with ASM. Before you configure the disks, you must determine the number of disks and the amount of free disk space that you require. The following sections describe how to identify the requirements and configure the disks on each platform:

- ["Identifying Storage Requirements for ASM"](#) on page 6-30

- ["Using an Existing ASM Disk Group"](#) on page 6-32
- ["Configuring Disks for ASM"](#) on page 6-33

Note: Although this section refers to disks, you can also use zero-padded files on a certified NAS storage device in an ASM disk group. See the appendices in *Oracle Database Installation Guide for UNIX Systems* for information about creating and configuring NAS-based files for use in an ASM disk group.

Identifying Storage Requirements for ASM

To identify the storage requirements for using ASM, you must determine how many devices and the amount of free disk space that you require. To complete this task, follow these steps:

1. Determine whether you want to use ASM for Oracle database files (datafiles), recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. One can use the file system, while the other uses ASM.

If you choose to enable automated backups and you do not have a cluster file system available, you must choose ASM for recovery file storage.

If you enable automated backups during the installation, you can choose ASM as the storage mechanism for recovery files by specifying an ASM disk group for the flash recovery area. Depending how you choose to create a database during the installation, you have the following options:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to use the same ASM disk group for datafiles and recovery files, or you can choose to use different disk groups for each file type.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must use the same ASM disk group for database files and recovery files.

2. Choose the ASM redundancy level that you want to use for the ASM disk group.

The redundancy level that you choose for the ASM disk group determines how ASM mirrors files in the disk group and determines the number of disks and amount of disk space that you require, as follows:

- 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 ASM does not mirror data in an external redundancy disk group, Oracle recommends that you use only RAID or similar devices that provide their own data protection mechanisms as disk devices in this type of disk group.

- Normal redundancy

In a normal redundancy disk group, ASM uses two-way mirroring by default, to increase performance and reliability. 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 most installations, Oracle recommends that you use normal redundancy disk groups.

- High redundancy

In a high redundancy disk group, 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.

While high redundancy disk groups do provide a high level of data protection, you must consider the higher cost of additional storage devices before deciding to use this redundancy level.

3. Determine the total amount of disk space that you require for the datafiles and recovery files.

Use the following table to determine the minimum number of disks and the minimum disk space requirements for the installation:

Redundancy Level	Minimum Number of Disks	Datafiles	Recovery Files	Both File Types
External	1	1.15 GB	2.3 GB	3.45 GB
Normal	2	2.3 GB	4.6 GB	6.9 GB
High	3	3.45 GB	6.9 GB	10.35 GB

You must also add additional disk space for the ASM metadata. You can use the following formula to calculate the additional disk space requirements (in MB):

$$15 + (2 * \text{number_of_disks}) + (126 * \text{number_of_ASM_instances})$$

For example, for a four-node RAC installation, using three disks in a high redundancy disk group, you require an additional 525 MB of disk space:

$$(15 + (2 * 3) + (126 * 4)) = 525$$

If an existing ASM instance exists on the system, 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 installation.

The following section describes how to identify existing disk groups and determine the free disk space that they contain.

4. Optionally identify failure groups for the ASM disk group devices.

Note: You need to complete this step only if you intend to use an installation method that runs DBCA in interactive mode, for example, if you intend to choose the Custom installation type or the Advanced database configuration option. Other installation types do not enable you to specify failure groups.

If you intend to use a normal or high redundancy disk group, 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 SCSI controller, the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To avoid failures of this type, you could use two SCSI controllers, each with two disks, and define a failure group for the disks attached to each controller. This configuration would enable the disk group to tolerate the failure of one SCSI controller.

Note: If you define custom failure groups, you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

5. If you are sure that a suitable disk group does not exist on the system, 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 ASM disk group should be the same size and have the same performance characteristics.
 - Do not specify more than one partition on a single physical disk as a disk group device. 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 ASM disk group, Oracle does not recommend their use. Logical volume managers can hide the physical disk architecture, preventing ASM from optimizing I/O across the physical devices.

For information about completing this task, see the "[Configuring Disks for ASM](#)" section on page 6-33.

Using an Existing ASM Disk Group

If you want to use ASM as the storage option for either database or recovery files, and an existing ASM disk group exists, you have the following choices, depending on the installation method that you select:

- If you select an installation method that runs DBCA in interactive mode, by choosing the Advanced database configuration option for example, you can decide whether you want to create a new disk group or use an existing one.

The same choice is available to you if you use DBCA after the installation to create a database.

- If you select an installation type that runs DBCA in non-interactive mode, you must choose an existing disk group for the new database; you cannot create a new disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The ASM instance that manages the existing disk group can be running in a different Oracle home directory.

To determine whether an existing ASM disk group exists, or to determine whether there is sufficient disk space in a disk group, you can use Oracle Enterprise Manager Database Grid Control or Database Control. Alternatively, you can use the following procedure:

1. View the contents of the `oratab` file to determine whether an ASM instance is configured on the system:

```
# more /var/opt/oracle/oratab
```

If an ASM instance is configured on the system, the `oratab` file should contain a line similar to the following:

```
+ASM:oracle_home_path:N
```

In this example, `+ASM` is the system identifier (SID) of the ASM instance and `oracle_home_path` is the Oracle home directory where it is installed. By convention, the SID for an ASM instance begins with a plus sign.

2. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the ASM instance that you want to use.
3. Connect to the ASM instance as the `SYS` user with `SYSDBA` privilege and start the instance if necessary:

```
# $ORACLE_HOME/bin/sqlplus "SYS/SYS_password as SYSDBA"
SQL> STARTUP
```

4. Enter the following command to view the existing disk groups, their redundancy level, and the amount of free disk space in each one:

```
SQL> SELECT NAME,TYPE,TOTAL_MB,FREE_MB FROM V$ASM_DISKGROUP;
```

5. From the output, identify a disk group with the appropriate redundancy level and note the free space that it contains.
6. If necessary, install or identify the additional disk devices required to meet the storage requirements listed in the previous section.

Note: If you are adding devices to an existing disk group, Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Configuring Disks for ASM

To configure disks for use with ASM, follow these steps:

1. If necessary, install the shared disks that you intend to use for the disk group and reboot the system.
2. To make sure that the disks are available, enter the following command:

```
# /usr/sbin/format
```

The output from this command is similar to the following:

```
AVAILABLE DISK SELECTIONS:
 0. c0t0d0 <ST34321A cyl 8892 alt 2 hd 15 sec 63>
    /pci@1f,0/pci@1,1/ide@3/dad@0,0
 1. c1t5d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
    /pci@1f,0/pci@1/scsi@1/sd@5,0
```

This command displays information about each disk attached to the system, including the device name (`cxydz`).

- From the list, identify the device names for the disk devices that you want to add to a disk group, then press `Ctrl/D` to exit from the `format` utility.
- Enter the following command to verify that the device you identified is not mounted as a file system:

```
# df -k
```

This command displays information about the partitions (slices) on disk devices that are mounted as file systems. The device name for a slice includes the disk device name, followed by the slice number, for example `cxydzsn`, where `sn` is the slice number. Slice 6 (`s6`) represents the entire disk. The disk devices that you choose must not be shown as mounted partitions.

- Enter the following commands to verify that the device you identified is not already part of an logical volume manager (LVM) disk group:

Note: The following command displays information about VERITAS Volume Manager disks. If you use a different LVM, see the appropriate documentation for information about determining which disk devices it is managing.

```
# vxdiskconfig
# /usr/sbin/vxdisk list
```

If this command displays disk group information associated with a disk device, the disk is already part of a disk group. The disks that you choose must not be part of a disk group.

- Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for each disk that you want to add to a disk group:

```
# chown oracle:dba /dev/rdisk/cxydzs6
# chmod 660 /dev/rdisk/cxydzs6
```

In this example, the device name specifies slice 6, which represents the entire disk.

Note: If you are using a multi-pathing disk driver with ASM, ensure that you set the permissions only on the correct logical device name for the disk.

Configure Raw Logical Volumes

See the following sections for information about configuring raw logical volumes or raw partitions:

- ["Create a Shared Disk Group"](#) on page 6-35
- ["Create Raw Logical Volumes in the New Disk Group"](#) on page 6-36
- ["Deport the Disk Group and Import It on the Other Cluster Nodes"](#) on page 6-38
- ["Create the DBCA Raw Device Mapping File"](#) on page 6-38

Note: If you are using ASM for database file storage, you need only create raw devices for the Oracle CRS files. However, Oracle recommends that you use the cluster file system to store the Oracle CRS files instead of using raw devices for them.

Create a Shared Disk Group

Note: This section describes how to configure raw logical volumes using VERITAS Cluster Volume Manager (CVM) with Sun Cluster 3.1. For information about other supported options for creating logical volumes for Oracle CRS and database files, see the Certify page on the Oracle *Metalink* Web site:

<http://metalink.oracle.com/>

To create a shared disk group, follow these steps:

1. If necessary, install the shared disks that you intend to use for the disk group and reboot the system.
2. To make sure that the disks are available, enter the following command:

```
# /usr/sbin/format
```

The output from this command is similar to the following:

```
AVAILABLE DISK SELECTIONS:
 0. c0t0d0 <ST34321A cyl 8892 alt 2 hd 15 sec 63>
    /pci@1f,0/pci@1,1/ide@3/dad@0,0
 1. c1t5d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
    /pci@1f,0/pci@1/scsi@1/sd@5,0
```

This command displays information about each disk attached to the system, including the device name (*cxydz*).

3. From the list, identify the device names for the disk devices that you want to add to a disk group, then press Ctrl/D to exit from the `format` utility.
4. Enter the following command to verify that the devices you identified are not mounted as file systems:

```
# df -k
```

This command displays information about the partitions (slices) on disk devices that are mounted as file systems. The device name for a slice includes the disk device name, followed by the slice number, for example *cxydzsn*, where *sn* is the slice number. Slice 2 (*s2*) represents the entire disk. The disk devices that you choose must not be shown as mounted partitions.

5. Enter the following commands to verify that the devices you identified are not already part of a disk group:

Note: The following command displays information about VERITAS Volume Manager disks. If you use a different LVM, see the appropriate documentation for information about determining which disk devices it is managing.

```
# /usr/sbin/vxdiskconfig
# /usr/sbin/vxdisk list
```

The `vxdisk list` command identifies the disk devices that are already configured in a disk group. The word `online` in the `STATUS` column also identifies disks that have been initialized and placed under VxVM control. The word `error` in the `STATUS` column identifies disks that are not initialized.

The disk devices that you choose must not be in an existing disk group.

- If the disk devices that you want to use are not initialized, enter a command similar to the following to initialize each disk:

```
# /usr/sbin/vxdiskadd cxydz
```

- To create a shared disk group, enter a command similar to the following, specifying all of the disks that you want to add to the group:

```
# /usr/sbin/vxdg -s init diskgroup diskname=devicename ...
```

In this example:

- `-s` indicates that you want to create a shared disk group
- `diskgroup` is the name of the disk group that you want to create, for example, `oradg`
- `diskname` is an administrative name that you assign to a disk, for example `orad01`
- `devicename` is the device name, for example, `c1t0d0`

Create Raw Logical Volumes in the New Disk Group

To create the required raw logical volumes in the new disk group, follow these steps:

- Choose a name for the database that you want to create.

The name that you choose must start with a letter and have no more than four characters, for example, `orcl`.

- From the following table, identify the logical volumes that you must create:

Number	Partition Size (MB)	Purpose and Sample Logical Volume Name
Raw Logical Volumes for Oracle Database		
1	500	SYSTEM tablespace: <code>dbname_system_raw_500m</code>
1	300 + (Number of instances * 250)	SYSAUX tablespace: <code>dbname_sysaux_raw_800m</code>
Number of instances	500	UNDOTBS n tablespace (One tablespace for each instance, where n is the number of the instance): <code>dbname_undotbsn_raw_500m</code>
1	160	EXAMPLE tablespace: <code>dbname_example_raw_160m</code>

Number	Partition Size (MB)	Purpose and Sample Logical Volume Name
1	120	USERS tablespace: <i>dbname_users_raw_120m</i>
2 * number of instances	120	Two online redo log files for each instance (where <i>n</i> is the number of the instance and <i>m</i> is the log number, 1 or 2): <i>dbname_redon_m_raw_120m</i>
2	110	First and second control files: <i>dbname_control[1 2]_raw_110m</i>
1	250	TEMP tablespace: <i>dbname_temp_raw_250m</i>
1	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwdfile_raw_5m</i>
Raw Logical Volumes for Oracle Cluster Ready Services (CRS)		
1	100	Oracle Cluster Registry: <i>ora_ocr_raw_100m</i> Note: You need to create this raw logical volume only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle cluster registry. If you are upgrading from Oracle9i Release 2, you can continue to use the raw device that you used for the SRVM configuration repository instead of creating this new logical volume.
1	20	Oracle clusterware voting disk: <i>ora_vote_raw_20m</i> Note: You need to create this raw logical volume only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle clusterware voting disk.

Note: The following examples show how to create a basic logical volume. See the VERITAS Volume Manager documentation for information about how to create more complex logical volumes using striping, mirroring, or both.

- To create the logical volume for the Oracle Cluster Registry, enter a command similar to the following:

```
# /usr/sbin/vxassist -g diskgroup make ora_ocr_raw_100m 100m user=root \
group=dba mode=660
```

In this example, *diskgroup* is the name of the disk group you created previously, for example, *oradg*.

- To create the other required logical volumes, enter a command similar to the following:

```
# /usr/sbin/vxassist -g diskgroup make volume size user=oracle \  
group=dba mode=660
```

In this example:

- *diskgroup* is the name of the disk group you created previously, for example, *oradg*.
- *volume* is the name of the logical volume that you want to create.
Oracle recommends that you use the sample names shown in the previous table for the logical volumes. Substitute the *dbname* variable in the sample logical volume name with the name you choose for the database in step 1.
- *size* is the size of the logical volume, for example, 500m represents 500 MB.
- `user=oracle group=dba mode=660` specifies the owner and permissions on the volume. Specify the Oracle software owner user and the OSDBA group for the user and group values (typically `oracle` and `dba`).

The following example shows a sample command used to create an 800 MB logical volume in the *oradg* disk group for the SYSAUX tablespace of a database named *test*:

```
# /usr/sbin/vxassist -g oradb make test_sysaux_800m 800m \  
user=oracle group=dba mode=660
```

Deport the Disk Group and Import It on the Other Cluster Nodes

To deport the disk group and import it on the other cluster nodes, follow these steps:

1. Deport the disk group:

```
# /usr/sbin/vxdg deport diskgroup
```

2. Log into each cluster node and complete the following steps:

- a. Enter the following command to cause VxVM to examine the disk configuration:

```
# /usr/sbin/vxdctl enable
```

- b. Import the shared disk group:

```
# /usr/sbin/vxdg -s import diskgroup
```

- c. Start all logical volumes:

```
# /usr/sbin/vxvol startall
```

Create the DBCA Raw Device Mapping File

Note: You must complete this procedure only if you are using raw devices for database files. You do not specify the raw devices for the Oracle CRS files in the DBCA raw device mapping file.

To enable Database Configuration Assistant (DBCA) to identify the appropriate raw device for each database file, you must create a raw device mapping file, as follows:

1. Set the ORACLE_BASE environment variable to specify the Oracle base directory that you identified or created previously:

- Bourne, Bash or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle ; export ORACLE_BASE
```

- C shell:

```
$ setenv ORACLE_BASE /u01/app/oracle
```

2. Create a database file subdirectory under the Oracle base directory and set the appropriate owner, group, and permissions on it:

```
# mkdir -p $ORACLE_BASE/oradata/dbname
# chown -R oracle:oinstall $ORACLE_BASE/oradata
# chmod -R 775 $ORACLE_BASE/oradata
```

In this example, *dbname* is the name of the database that you chose previously.

3. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
4. Enter the following command to create a text file that you can use to create the raw device mapping file:

```
# find /dev/vx/rdisk/diskgroup -user oracle -name dbname* \
-print > dbname_raw.conf
```

5. Edit the `dbname_raw.conf` file in any text editor to create a file similar to the following:

Note: The following example shows a sample mapping file for a two-instance RAC cluster.

```
system=/dev/vx/rdsk/diskgroup/dbname_system_raw_500m
sysaux=/dev/vx/rdsk/diskgroup/dbname_sysaux_raw_800m
example=/dev/vx/rdsk/diskgroup/dbname_example_raw_160m
users=/dev/vx/rdsk/diskgroup/dbname_users_raw_120m
temp=/dev/vx/rdsk/diskgroup/dbname_temp_raw_250m
undotbs1=/dev/vx/rdsk/diskgroup/dbname_undotbs1_raw_500m
undotbs2=/dev/vx/rdsk/diskgroup/dbname_undotbs2_raw_500m
redo1_1=/dev/vx/rdsk/diskgroup/dbname_redo1_1_raw_120m
redo1_2=/dev/vx/rdsk/diskgroup/dbname_redo1_2_raw_120m
redo2_1=/dev/vx/rdsk/diskgroup/dbname_redo2_1_raw_120m
redo2_2=/dev/vx/rdsk/diskgroup/dbname_redo2_2_raw_120m
control1=/dev/vx/rdsk/diskgroup/dbname_control1_raw_110m
control2=/dev/vx/rdsk/diskgroup/dbname_control2_raw_110m
spfile=/dev/vx/rdsk/diskgroup/dbname_spfile_raw_5m
pwdfile=/dev/vx/rdsk/diskgroup/dbname_pwdfile_raw_5m
```

In this example:

- *diskgroup* is the name of the volume group
- *dbname* is the name of the database

Use the following guidelines when creating or editing this file:

- Each line in the file must have the following format:

```
database_object_identifier=logical_volume
```

The logical volume names suggested in this manual include the database object identifier that you must use in this mapping file. For example, in the following logical volume name, `redo1_1` is the database object identifier:

```
/dev/vx/rdsk/oradg/rac_redo1_1_raw_120m
```

- For a single instance database, the file must specify one automatic undo tablespace datafile (`undotbs1`), and at least two redo log files (`redo1_1`, `redo1_2`).
 - For a RAC database, the file must specify one automatic undo tablespace datafile (`undotbsn`) and two redo log files (`redon_1`, `redon_2`) for each instance.
 - Specify at least two control files (`control1`, `control2`).
 - To use manual instead of automatic undo management, specify a single RBS tablespace datafile (`rbs`) instead of the automatic undo management tablespaces.
6. Save the file and note the file name that you specified.
 7. When you are configuring the `oracle` user's environment later in this chapter, set the `DBCA_RAW_CONFIG` environment variable to specify the full path to this file.

Verify that the Cluster Software is Configured and Running

Note: The cluster software mentioned in this section is not required to install Oracle Real Application Clusters. However, if this software is installed, Oracle Cluster Ready Services (CRS) can integrate with it.

To verify that the Sun Cluster software is configured and running:

1. Verify that the `oracle.dba.gid` parameter in the `/etc/opt/SUNWcluster/conf/clustername.ora_cdb` file specifies the group ID of the OSDBA group (typically `dba`) that you created earlier.

Note: The name of the cluster (`clustername`) is defined in the `/etc/opt/SUNWcluster/conf/default_clustername` file.

2. To verify that the Sun Cluster software is running, enter the following command on one of the cluster nodes:

```
# /usr/cluster/bin/scstat -n
```

3. If the output from this command does not list all of the nodes in the cluster, make sure that all of the cluster nodes are booted.

If necessary, see the Sun Cluster documentation for information about troubleshooting the cluster.

Stop Existing Oracle Processes

Caution: If you are installing additional Oracle Database 10g products in an existing Oracle home, stop **all** processes running in the Oracle home. You must complete this task to enable the Installer to relink certain executables and libraries.

If you choose to create a database during the installation, most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the IPC key value EXTPROC. However, if an existing Oracle Net listener process is using the same port or key value, the Installer can only configure the new listener; it cannot start it. To ensure that the new listener process starts during the installation, you must shut down any existing listeners before starting the Installer.

To determine whether an existing listener process is running and to shut it down if necessary, follow these steps:

1. Switch user to `oracle`:

```
# su - oracle
```

2. Enter the following command to determine whether a listener process is running and to identify its name and the Oracle home directory in which it is installed:

```
$ ps -ef | grep tnslnsr
```

This command displays information about the Oracle Net listeners running on the system:

```
... oracle_home1/bin/tnslnsr LISTENER -inherit
```

In this example, `oracle_home1` is the Oracle home directory where the listener is installed and `LISTENER` is the listener name.

3. Set the `ORACLE_HOME` environment variable to specify the appropriate Oracle home directory for the listener:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=oracle_home1
$ export ORACLE_HOME
```

- C or tcsh shell:

```
$ setenv ORACLE_HOME oracle_home1
```

4. Enter the following command to identify the TCP/IP port number and IPC key value that the listener is using:

```
$ $ORACLE_HOME/bin/lsnrctl status listenername
```

Note: If the listener uses the default name `LISTENER`, you do not have to specify the listener name in this command.

5. Enter a command similar to the following to stop the listener process:

```
$ $ORACLE_HOME/bin/lsnrctl stop listenername
```

6. Repeat this procedure to stop all listeners running on this system and on all other cluster nodes.

Configure the oracle User's Environment

You run the Installer from the `oracle` account. However, before you start the Installer you must configure the environment of the `oracle` user. To configure the environment, you must:

- Set the default file mode creation mask (`umask`) to 022 in the shell startup file.
- Set the `DISPLAY` and `ORACLE_BASE` environment variables.

To set the `oracle` user's environment, follow these steps:

1. Start a new terminal session, for example, an X terminal (`xterm`).
2. Enter the following command to ensure that X Window applications can display on this system:

```
$ xhost +
```

3. If you are not already logged in to the system where you want to install the software, log in to that system as the `oracle` user.
4. If you are not logged in as the `oracle` user, switch user to `oracle`:

```
$ su - oracle
```

5. To determine the default shell for the `oracle` user, enter the following command:

```
$ echo $SHELL
```

6. Open the `oracle` user's shell startup file in any text editor:

- Bourne shell (`sh`), Bash shell (`bash`), or Korn shell (`ksh`):

```
$ vi .profile
```

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

```
% vi .login
```

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

```
umask 022
```

8. If the `ORACLE_SID`, `ORACLE_HOME`, or `ORACLE_BASE` environment variables are set in the file, remove the appropriate lines from the file.

9. Save the file and exit from the editor.

10. To run shell startup script, enter one of the following command:

- Bourne, Bash, or Korn shell:

```
$ . ./profile
```

- C shell:

```
% source ./login
```

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

- Bourne, Bash, or Korn shell:

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

- C shell:

```
$ setenv DISPLAY local_host:0.0
```

In this example, *local_host* is the host name or IP address of the system you want to use to display the Installer (your workstation or PC).

12. If you determined that the `/tmp` directory has less than 200 MB of free disk space, identify a file system with at least 200 MB of free space and set the `TEMP` and `TMPDIR` environment variables to specify a temporary directory on this file system:

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

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

```
$ su - root
# mkdir /mount_point/tmp
# chmod a+wr /mount_point/tmp
# exit
```

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

- * Bourne, Bash, or Korn shell:

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

- * C shell:

```
$ setenv TEMP /mount_point/tmp
$ setenv TMPDIR /mount_point/tmp
```

13. Enter commands similar to the following to set the `ORACLE_BASE` environment variable:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_BASE=/u01/app/oracle
$ export ORACLE_BASE
```

- C shell:

```
% setenv ORACLE_BASE /u01/app/oracle
```

In these examples, `/u01/app/oracle` is the Oracle base directory that you created earlier.

14. If necessary, set the following environment variables. If you must specify several values for an environment variable, for example `PATH`, separate the values with a colon (`:`).

Environment Variable	Required By	Sample Setting and Description
DBCA_RAW_CONFIG	Oracle Database installations that use raw devices for database storage.	\$ORACLE_BASE/oradata/dbname/dbname_raw.conf Specifies the location of the raw device mapping file
COBDIR	Pro*COBOL	/opt/lib/cobol Specifies the directory where COBOL is installed on the system. Note: To install Pro*COBOL, you must choose a custom software installation.
PATH	Pro*COBOL	\$COBDIR/bin:\$PATH Specifies the directory where the COBOL compiler executable is located.

15. Enter the following command to ensure that the ORACLE_HOME and TNS_ADMIN environment variable are not set:

```
$ unset ORACLE_HOME
$ unset TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value it specifies as the default path for the Oracle home directory. However, if you set the ORACLE_BASE environment variable, Oracle recommends that you unset this environment variable and choose the default path suggested by the Installer.

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

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

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

Part III

Installing CRS and Oracle Database 10g with RAC, Creating RAC Databases, and Performing Post-Installation Tasks

Part III describes the two-phase installation process of how to install Cluster Ready Services (CRS) and the Oracle Database 10g with Real Application Clusters (RAC). It also explains how to create RAC databases, and it describes the post-installation tasks. The chapters in Part III are:

- [Chapter 7, "Installing Cluster Ready Services"](#)
- [Chapter 8, "Installing Oracle Database 10g with Real Application Clusters"](#)
- [Chapter 9, "Creating RAC Databases with the Database Configuration Assistant"](#)
- [Chapter 10, "Real Application Clusters Post-Installation Procedures"](#)

Installing Cluster Ready Services

This chapter describes the installation procedures for phase one for installing Cluster Ready Services for Oracle Database 10g Real Application Clusters. This chapter contains the following topics:

- [Installation Setup Procedures](#)
- [Installing the UDLM for SunClusters](#)
- [Installing Cluster Ready Services with the OUI](#)

Installation Setup Procedures

Perform the following procedures to complete phase one of the installation of the Oracle Database 10g with RAC.

1. Verify user equivalence by executing the `ssh` command on the local node with the `date` command argument using the following syntax:

```
ssh node_name date
```

The output from this command should be the timestamp of the remote node identified by the value that you use for `node_name`. If `ssh` is in the `/usr/local/bin` directory, then use `ssh` to configure user equivalence.

You cannot use `ssh` to verify user equivalence if `ssh` is in another location in your `PATH`. In this case, use `rsh` to confirm user equivalence.

Note: When you test user equivalence by executing the `ssh` or `rsh` commands, the system should not respond with questions nor should you see additional output, for example, besides the output of the `date` command.

2. In addition to the host machine's public internet protocol (IP) address, obtain two more IP addresses for each node that is going to be part of your installation. During the installation, enter the IP addresses into DNS. One of the IP addresses must be a *public* IP address for the node's virtual IP address (VIP). Oracle uses VIPs for client-to-database connections. Therefore, the VIP address must be publicly accessible. The other address must be a *private* IP address for intra-node, or instance-to-instance Cache Fusion traffic. Using public interfaces for Cache Fusion can cause performance problems.

Installing the UDLM for SunClusters

If you are using Sun Clusters, then you must install the Oracle-provided UDLM patch onto each node that is part of your current cluster installation. You must install the UDLM patch *before* you install Cluster Ready Services. Even if you have a pre-Oracle Database 10g UDLM, you *must* install the Oracle Database 10g UDLM.

Install the UDLM patch using the procedures in the `README` file that is located in the `/racpatch` directory on the Oracle Cluster Ready Services Release 1 (10.1.0.2) CD-ROM.

Note: You must stop the cluster and upgrade the UDLM one node at a time if you are upgrading from a previous UDLM version.

Installing Cluster Ready Services with the OUI

This section describes the procedures for using the Oracle Universal Installer (OUI) to install CRS. Note that the CRS home that you identify in this phase of the installation is only for CRS software; this home *cannot* be the same home as the home that you will use in phase two to install the Oracle Database 10g software with RAC.

Note: You cannot install Oracle Database 10g Cluster Ready Services software on an Oracle Database 10g cluster file system.

1. If you are installing CRS on a node that already has a single-instance Oracle Database 10g installation, stop the existing ASM instances and Cluster Synchronization Services (CSS) daemon and use the script `$ORACLE_HOME/bin/localconfig delete` in the home that is running CSS to reset the OCR configuration information. After CRS is installed, then start up the ASM instances again and the ASM instances will use the cluster CSS daemon instead of the daemon for the single-instance Oracle database.
2. Run the `runInstaller` command from the `/crs` subdirectory on the Oracle Cluster Ready Services Release 1 (10.1.0.2) CD-ROM. This is a separate CD that contains the Cluster Ready Services software. When the OUI displays the Welcome page, click Next.
3. Depending on whether your environment has an Oracle inventory, the following scenarios apply:
 - If you are performing this installation in an environment where the OUI inventory is already set up, then the OUI displays the Specify File Locations page. If the Specify File Locations page appears, proceed to Step 5.
 - If you are performing this installation in an environment in which you have never installed Oracle database software, in other words the environment does not have an OUI inventory, then the OUI displays the Specify Inventory Directory and Credentials page. Enter the inventory location and the UNIX group name information into the Specify Inventory Directory and Credentials page, click Next, and the OUI displays a dialog.
4. The OUI dialog indicates that you should run the `/oraInstRoot.shoraInventory location` script. Run the `oraInstRoot.sh` script as root user, click Continue, and the OUI displays the Specify File Locations page.
5. The Specify File Locations Page contains predetermined information for the source of the installation files and the target destination information. Enter the CRS home name and its location in the target destination, click Next, and the OUI displays the Language Selection page.

Note: The CRS home that you identify in this step *must be different* from the Oracle home that you will use in phase two of the installation.

6. In the Language Selection page select the languages that you want CRS to use, click Next, and the OUI displays the Cluster Configuration Information page.
7. The Cluster Configuration Information page contains pre-defined node information if the OUI detects that your system has vendor clusterware. Otherwise, the OUI displays the Cluster Configuration Information page without pre-defined node information.

If you install the clusterware in this installation session without using vendor clusterware, then enter a public node name and a private node name for each node. 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. This node name can be either the permanent or the virtual host name.

In addition, the cluster name that you use must be globally unique throughout the enterprise and the allowable character set for the cluster name is the same as that of hostnames. For example, you cannot use the characters `!@#%^&* or ()`. Oracle recommends that you use the vendor cluster name if one exists. Make sure

that you also enter a private node name or private IP address for each node. This is an address that is only accessible by the other nodes in this cluster. Oracle uses the private IP addresses for Cache Fusion processing. Click Next after you have entered the cluster configuration information and the OUI performs validation checks such as node availability and remote Oracle home permissions verifications. These verifications may require some time to complete. When the OUI completes the verifications, the OUI displays the Private Interconnect Enforcement page.

Note: The IP addresses that you use for all of the nodes in the current installation process must be from the same subnet.

8. In the Private Interconnect Enforcement page the OUI displays a list of cluster-wide interfaces. Use the drop-down menus on this page to classify each interface as `Public`, `Private`, or `Do Not Use`. The default setting for each interface is `Do Not Use`. You *must* classify at least one interconnect as `Public` and one as `Private`.
9. When you click Next on the Private Interconnect Enforcement page, if the `ocr.loc` file exists in the `/etc` directory (In Solaris Operating System (SPARC 64-bit) and HP-UX environments, the `ocr.loc` file is in the `/var/opt/oracle` directory), and if the `ocr.loc` file has a valid entry for the Oracle Cluster Registry (OCR) location, then the Voting Disk Location page appears and you should proceed to Step 10.

Otherwise, the Oracle Cluster Registry Location Information page appears. Enter a complete path for the raw device or shared file system file for the Oracle Cluster Registry, click Next, and the Voting Disk Information page appears.

10. On the Voting Disk Information Page, enter a complete path and file name for the file in which you want to store the voting disk and click Next. This must be a shared raw device or a shared file system file.

Note: The storage size for the OCR should be at least 100MB and the storage size for the voting disk should be at least 20MB. In addition, Oracle recommends that you use a RAID array for storing the OCR and the voting disk to ensure the continuous availability of the partitions.

See Also: The pre-installation chapters in [Part II](#) for information about the minimum raw device sizes

11. After you complete the Voting Disk Information page and click Next, if the Oracle inventories on the remote nodes are not set up, then the OUI displays a dialog asking you to run the `oraInstRoot.sh` script on all of the nodes. After the `oraInstRoot.sh` script processing completes, the OUI displays a Summary page.
12. Verify that the OUI should install the components shown on the Summary page and click Install.

During the installation, the OUI first copies software to the local node and then copies the software to the remote nodes. Then the OUI displays a dialog indicating that you must run the `root.sh` script on all the nodes that are part of this installation. Execute the `root.sh` script on *one node at a time* and click OK in the

dialog that `root . sh` displays after it completes each session. Only start another session of `root . sh` on another node after the previous `root . sh` execution completes; do *not* execute `root . sh` on more than one node at a time. When you complete the final execution of `root . sh`, the `root . sh` script runs the following assistants without your intervention:

- Oracle Cluster Registry Configuration Tool (`ocrconfig`)—If this tool detects a 9.2.0.2 version of RAC, then the tool upgrades the 9.2.0.2 OCR block format to an Oracle Database 10g OCR block format.
- Cluster Configuration Tool (`clscfg`)—This tool automatically configures your cluster and creates the OCR keys.

13. When the OUI displays the End of Installation Page, click Exit to exit the Installer.

14. Verify your CRS installation by executing the `olsnodes` command from the `CRS Home/bin` directory. The `olsnodes` command syntax is:

```
olsnodes [-n] [-l] [-v] [-g]
```

Where:

- `-n` displays the member number with the member name
- `-l` displays the local node name
- `-v` activates verbose mode
- `-g` activates logging

The output from this command should be a listing of the nodes on which CRS was installed.

15. After you run `root . sh` on all of the nodes and click OK on the last `root . sh` dialog, the OUI runs the Oracle Notification Server Configuration Assistant and Oracle Private Interconnect Configuration Assistant. These assistants run without user intervention.

At this point, you have completed phase one, the installation of Cluster Ready Services, and are ready to install the Oracle Database 10g with RAC as described in [Chapter 8, "Installing Oracle Database 10g with Real Application Clusters"](#).

Cluster Ready Services Background Processes

The following processes must be running in your environment after the CRS installation in order for Cluster Ready Services to function:

- `oproc` -- Process monitor for the cluster.
- `evmd` -- Event manager daemon that starts the `racgevt` process to manage callouts.
- `ocssd` -- Manages cluster node membership and runs as `oracle` user; failure of this process results in cluster restart.
- `crsd` -- Performs high availability recovery and management operations such as maintaining the OCR. Also manages application resources and runs as `root` user and restarts automatically upon failure.

Installing Oracle Database 10g with Real Application Clusters

This chapter describes the phase two installation procedures for installing the Oracle Database 10g with Real Application Clusters (RAC). This chapter also describes some of the Oracle Universal Installer (OUI) features. The topics in this chapter are:

- [Selecting a Database Configuration Type](#)
- [Installation Setup Procedures](#)
- [Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer](#)

Selecting a Database Configuration Type

This section describes OUI features that you should understand before beginning phase two of the RAC installation process. When you run the OUI and select the Oracle Database 10g, you can select the General Purpose, Transaction Processing, Data Warehouse, or Customized database configuration type.

For the first three configuration types, you can complete additional procedures that are described later in this chapter. If you select the fourth type, or the Customized configuration, then you can use the DBCA to create the database as described in [Chapter 9](#). Oracle recommends that you use the DBCA to create your database.

You can also select the Customized configuration, select a preconfigured template, customize the template, and use the DBCA to create a database using the template. These templates correspond to the General Purpose, Transaction Processing, and Data Warehouse configuration types. You can also use the DBCA with the Custom template to create a database.

To manually configure your environment and manually create your database, select the Software Only database configuration option and refer to the manual database creation procedures posted at <http://otn.oracle.com>. However, Oracle recommends that you use one of the preconfigured database options or use the Custom option and the DBCA to create your database.

Configuration Type Descriptions

The configuration type that you select, as described in [Table 8–1](#), determines how you proceed.

Table 8–1 Oracle Universal Installer Database Configuration Types

Configuration Type	Description	Advantages
General Purpose, Transaction Processing, and Data Warehouse	Installs a preconfigured starter database, licensable Oracle options (including the Oracle Database 10g with RAC), networking services, Oracle Database 10g utilities, and online documentation. At the end of the installation, the DBCA creates and configured your RAC database.	Minimal input required. You can create your database more quickly than with the Customized type.
Customized	Enables you to customize your database options and storage components.	Enables you to create arbitrary tablespaces and datafiles and customize all aspects of your database.
Software Only	Installs only the software. Does not configure the listeners or network infrastructure and does not create a database.	

General Purpose, Transaction Processing, and Data Warehouse Configuration Types

If you select one of the first three configuration types on the OUI Database Configuration page, then complete one of the platform-specific procedures in this chapter under the heading "[Installation Setup Procedures](#)" on page 8-4. These three configuration types use preconfigured templates. After you complete these procedures, the Oracle Net Configuration Assistant and the DBCA run without further input and the OUI displays a progress indicator during the installation.

The DBCA processing for these configuration types creates a starter database and configures the Oracle network services. If you do choose raw devices on the Storage Options page, then the DBCA verifies that you configured the raw devices for each tablespace.

Note: If you have not properly configured the raw devices as described in [Appendix C](#), then the DBCA cannot create your database.

If you select the Customized configuration, then you must enter specific information as described in the next section.

Using The Customized Configuration Type

When you select the Customized configuration type, the OUI runs the DBCA which displays four preconfigured database template choices:

- General Purpose
- Transaction Processing
- Data Warehouse
- Custom

The first three templates create a database that is optimized for that environment. You also can customize these templates. The Custom type, however, creates a database without using preconfigured options. Therefore, if you are not using a cluster file system, then you *must* enter specific raw device information for the datafiles to perform a Custom installation as described in [Chapter 9, "Creating RAC Databases with the Database Configuration Assistant"](#).

The following section provides more detail about OUI and DBCA processing when creating a RAC database.

Processing of the OUI, the DBCA, and Other Assistants During Installation

After installation, the OUI starts the Oracle Net Configuration Assistant. After the Oracle Net Configuration Assistant completes its processing, the OUI runs the DBCA to create your database using the optimal flexible architecture (OFA). This means that the DBCA creates your database files, including the default server parameter file (spfile), using standard file naming and file placement practices. The primary phases of DBCA processing are:

- Verify that you correctly configured the shared disks for each tablespace if you are using raw storage
- Create the database
- Configure the Oracle network services
- Start the listeners and database instances

You can also use the DBCA in standalone mode to create a database.

See Also: The *Oracle Net Services Administrator's Guide* if you experience problems, for example, with the listener configuration, and for further information about LDAP support

The remainder of this chapter explains how to use the OUI to install the Oracle Database 10g with RAC.

Installation Setup Procedures

If the user who installed CRS is different from the user who installs RAC, then all of the user-level pre-install steps must be completed.

1. Verify user equivalence by executing the `ssh` command on the local node with the `date` command argument using the following syntax:

```
ssh node_name date
ssh node_name xclock
```

The output from these commands should be the timestamp of the remote node identified by the value that you use for `node_name`. In addition, the system should display the remote node's `xclock`. When you execute these commands, you should not see any other errors, warnings, or additional output. If `ssh` is in the `/usr/local/bin` directory, then use `ssh` to configure user equivalence.

The OUI cannot use `ssh` to verify user equivalence if `ssh` is in another location in your `PATH`. In this case, use `rsh` to confirm user equivalence.

Note: When you test user equivalence by executing the `ssh` or `rsh` commands, the system should not respond with questions nor should you see additional output, for example, besides the output of the `date` command and the `xclock`. You should also not have any echo messages in your `.login` or `.cshrc` files.

2. Perform a network connectivity test to ensure that all nodes on which you want to perform this installation can communicate with each other. Ensure that the public and private network interfaces have the same interface names on each node of your cluster.
3. Set directory and file creation permissions using the `umask` command so that you can write as `oracle` user to all of the Oracle homes in the RAC database environment in which you are performing this installation.
4. Create directories for the Oracle home and Oracle datafiles according to OFA standards.

Note: The Oracle home that you create for installing Oracle Database 10g with the RAC software cannot be the same Oracle home that you used during the clusterware installation.

Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer

Perform the following procedures to install the Oracle Database 10g with RAC.

1. Execute the `runInstaller` command from the `root` directory of the Oracle Database 10g Release 1 (10.1.0.2) CD-ROM.
2. When the OUI displays the Welcome page, click Next, and the OUI displays the Specify File Locations page.
3. The Source field on the Specify File Locations page is pre-populated with the path to the file Oracle Database 10g `products.xml`. Enter the Oracle home name and location in the target destination and click Next.

Note: The Oracle home that you use in this step *must be different* from the home that you used during the CRS installation in phase one. In other words, you **must not** install Oracle Database 10g with RAC software into the same home in which you installed the CRS software.

If you enter a pre-existing Oracle home and the Oracle home is not registered with the OUI inventory, then the OUI displays a warning that the Oracle home that you selected is non-empty. Click OK on this OUI warning dialog to continue installing on that directory. If the Oracle home exists and it is registered with the OUI inventory and the Oracle has a cluster home, then the OUI displays the Selected Nodes page as described in Step 4. If you enter an Oracle home that does not exist, then the OUI displays the Specify Hardware Cluster Installation Mode page described in Step 5 of this procedure.

Note: If the OUI does not detect an inventory entry for a CRS installation, or if the Oracle home that you select is an existing non-RAC, single-node home, then the OUI assumes that you are performing a single-node, non-RAC installation.

4. The Selected Nodes page is an informational page that displays the selected nodes that are associated with a cluster home. Click Next on this page and the OUI displays the Install Type page described in Step 6.

When you click Next on the Selected Nodes page, the OUI verifies that the Oracle home directory is writable on the remote nodes and that the remote nodes are operating. The OUI also re-validates user equivalence.

If the OUI detects a network or user equivalence problem on any node that you have included in this installation, then the OUI displays a warning on the Selected Nodes page. This warning appears next to the node and indicates that you should correct a problem on the affected node before proceeding. To resolve problems, examine the OUI actions recorded in the following file:

```
OraInventory/logs/installActionsdate_time.log
```

5. On the Specify Hardware Cluster Installation Mode page, select an installation mode. The Cluster Installation mode is selected by default when the OUI detects that you are performing this installation on a cluster. In addition, the local node is always selected for the installation. Select additional nodes that are to be part of this installation session and click Next.

See Also: If you select Local Installation, then refer to the *Oracle Database Installation Guide, 10g Release 1 (10.1) for UNIX Systems* to perform a single-node, non-RAC installation on a cluster, or refer to the appropriate platform-specific quick start guide

When you click Next on the Specify Hardware Cluster Installation page, the OUI verifies that the Oracle home directory is writable on the remote nodes and that the remote nodes are operating. The OUI also re-validates user equivalence.

If the OUI detects a network or user equivalence problem on any node that you have included in this installation, then the OUI displays a warning on the Specify Hardware Cluster Installation Selection page. This warning appears next to the node and indicates that you should correct a problem on the affected node before

proceeding. To resolve problems, examine the OUI actions recorded in the following file:

```
OraInventory/logs/installActionsdate_time.log
```

You can either make corrections and click Next, or you can deselect nodes that have errors and click Next and the OUI displays the Install Type page.

Note: Each node in a cluster must have a CRS home to enable the clusterware to operate on each node. However, the Oracle home can be on a subset of the nodes that are part of your installation.

6. On the Install Type page you can select Enterprise Edition, Standard Edition, or Custom Install type. If you select Enterprise Edition or Standard Edition and click Next, then the OUI displays the Prerequisite Check page. The Custom install type enables you to select the individual components that you want to install from the list of available components. If you have purchased a Standard Edition license and you perform a Custom installation, then ensure that you only install the products covered by the Standard Edition license.
7. The Prerequisite Check page verifies the operating system kernel parameters or attributes and calculates the `ORACLE_BASE` location. Click Next. If the OUI discovers a database to upgrade, then the OUI displays the Upgrading an Existing Database page that lists the databases that are available to upgrade. Select a database to upgrade as described in Step 8. If you do not select a database to upgrade and click Next, then the OUI displays the Create Starter Database page as described in Step 9. Additionally, if the OUI does not find any databases to upgrade, then the OUI displays the Create Starter Database Page as described in Step 9 if you choose Enterprise Edition or Standard Edition.
8. On the Upgrading an Existing Database page, select the check box labelled "Upgrade an Existing Database" and then select the databases to upgrade from the list below. Click Next, and this by-passes the database creation process and the OUI displays the Summary page as described in Step 15.
9. On the Create a Starter Database Page select General Purpose, Transaction Processing, Data Warehouse, or Advanced. Choose to either create a database or to perform a Software Only installation, click Next, and the OUI displays the Database Configuration Options Page.

If you select "Do Not Create Starter Database", then select a Software Only installation. Click Next, and the OUI performs a software installation without performing any configuration tasks.

If you choose the "Custom Install Type", then the OUI performs a prerequisite check for the operating system kernel parameters or attributes and the OUI calculates the `ORACLE_BASE` location.

After you make your selection on this page and click Next, the OUI displays the Database Configuration Options page.

10. On the Database Configuration Options page enter a global database name. A global database name is a name that includes the database name and database domain, such as `db.us.acme.com`. The name that you enter on this page must be unique among all the global database names used in your environment. Accept or change the common prefix for the Oracle `sid` for each instance. Each instance has a `sid` that comprises the common prefix that you enter in this step and an instance ID that is automatically generated. Also select a database character set

and select any database examples to install. Note that a SID cannot exceed five characters in length. Click Next and the OUI displays the Database Management Option for Database Control page (Grid or Local Database). If you choose the local Database Control, then you can choose the email option and enter the outgoing SMTP server name and email address. If there is an existing Enterprise Manager Grid available, then there will be a drop-down menu in which you can indicate that you want to use Enterprise Manager Grid Control for database management. Click Next, and the OUI displays the Data Storage Option page.

11. Select a data storage option on the Data Storage Option Page.

If you select Automatic Storage Management (ASM)—Then the OUI displays an ASM Management Options page showing the disk partition locations from which you must select. These partitions appear in the default location which is in following format:

```
/dev/rdsk/*
```

If you select File System—Then enter the full path of the location for the datafiles destination on your shared or cluster file system and click Next.

If you select raw devices—Then enter a location for the raw device mapping file and click Next. If the `DBCA_RAW_CONFIG` environment variable is set, then the location already is set with the value for this variable.

When you click Next on the Data Storage Option page, the OUI displays the Backup and Recovery page.

12. If you enable backup on the Backup and Recovery page, then you can choose File System or ASM. You will also enter the user name and password. When you click Next on the Backup and Recovery page, the OUI displays the Schema Password page.
13. On the Schema Password page you can choose different passwords for SYS, SYSTEM, DBSNMP, and SYSMAN, or you can choose one password for all of the privileged accounts. When you click Next on the Schema Password page, the OUI displays the Operation System Group page.
14. On the Operation System Group page enter the group name for the SYSDBA and SYSOPER users. When you click Next on the Operation System Group page, the OUI displays the Summary page.
15. The Summary Page displays the software components that the OUI will install and the space available in the Oracle home with a list of the nodes that are part of the installation session. Verify the details about the installation that appear on the Summary page and click Install or click Back to revise your installation.

During the installation, the OUI copies software to the local node and then copies the software to the remote nodes. Then the OUI prompts you to run the `root.sh` script on all the selected nodes. The `root.sh` script starts the Virtual Internet Protocol Configuration Assistant (VIPCA) and displays the VIPCA Welcome page. Before you run `root.sh`, make sure that your display environment variables are properly set.

16. Review the information on the VIPCA Welcome page, click Next, and the VIPCA displays the Public Network Interfaces page.
17. On the Public Network Interfaces page determine the network interface cards (NICs) to which you want to assign your public VIP addresses, click Next, and the VIPCA displays the IP Address page.

18. On the IP Address page enter an *unused* (unassigned) public virtual IP address for each node displayed on this page and click Next. After you click Next, the VIPCA displays a Summary page. Review the information on this page and click Finish. A progress dialog appears while the VIPCA configures the virtual IP addresses with the network interfaces that you selected. The VIPCA then creates and starts the VIPs, GSD, and Oracle Notification Service (ONS) node applications. When the configuration completes, click OK to see the VIPCA session results. Review the information on the Configuration Results page, and click Exit to exit the VIPCA.
19. Repeat the `root . sh` procedure on all nodes that are part of this installation. The VIPCA will not run again on the remote node because the remote node is already configured.
20. After running `root . sh` on all the nodes in the previous step, click OK on the OUI dialog to continue the installation. This enables the remaining Oracle configuration assistants to run in the order shown in the following list so that the assistants can perform post-installation processing:
 - Net Configuration Assistant (NetCA)
 - Database Configuration Assistant (DBCA)

At the end of the installation, the OUI starts Enterprise Manager Database Control in standalone mode. Database Control is immediately available to administer and monitor your newly installed Oracle Database 10g with Real Application Clusters environment.

You have completed the second and final phase of the installation. Proceed to [Chapter 10, "Real Application Clusters Post-Installation Procedures"](#) to perform the post-installation tasks.

Note: Use the DBCA and the OUI to deinstall RAC.

De-Installing Real Application Clusters Software

Execute the following procedures to de-install Oracle Database 10g RAC software. First deinstall the Oracle database software, then deinstall the Cluster Ready Services (CRS) software as described in the following sections:

- [De-Installing Oracle Database 10g RAC Software](#)
- [De-Installing Cluster Ready Services](#)

De-Installing Oracle Database 10g RAC Software

Execute the following procedures to deinstall the Oracle Database 10g RAC software:

1. Remove the Oracle Database 10g with RAC by using the DBCA to drop the database.
2. If this is the Oracle home from which the per-node listener named `LISTENER_nodename` runs, then use NetCA to remove this listener and its CRS resources. If necessary, re-create this listener in another home.
3. If this is the Oracle home from which the per-node ASM instance runs, then use the command `srvctl stop asm -n node` followed by `srvctl remove asm -n node` to remove its CRS resources. Remove any `oratab` entry beginning with `+ASM`. If necessary, re-create the ASM instance in another home.

4. If VIP, ONS and GSD node applications are created using this Oracle home, then you must stop and remove these node applications. To stop CRS node applications, execute the command `srvctl stop nodeapps -n node`. Repeat the same command for each node in the cluster. To remove CRS node applications, execute the command `$ORACLE_HOME/install/rootdeletenode.sh`. If other Oracle Database 10g RAC Oracle homes exist, then recreate the node applications by running the `root.sh` script in *that* Oracle home.

See Also: *Oracle Real Application Clusters Administrator's Guide* for more information about using RAC scalability features of adding and deleting nodes and instances from RAC databases

Note: You cannot perform a RAC installation from the same OUI session in which you perform a RAC de-installation. In other words, if you deinstall RAC with the OUI and want to perform another RAC installation, then you must start a new OUI session.

De-Installing Cluster Ready Services

Execute the following procedures to deinstall the CRS software:

1. Deinstall each Oracle Database 10g RAC home by executing the procedures under the previous heading, "[De-Installing Oracle Database 10g RAC Software](#)".
2. Execute the command `CRSHome/install/rootdelete.sh` to disable the CRS applications that are on running on the cluster node. The `rootdelete.sh` script requires three arguments. If you are running this command on a remote node of the cluster then use `remote`, otherwise use `local` as the first argument. If the `ocr.loc` file is on a shared file system, then use `sharedvar`. Otherwise use `nosharedvar` as the second argument. If the CRS home is on a shared file system, then use `sharedhome`, otherwise use `nosharedhome` as the third argument. Repeat this step for each node of the cluster from which you want to deinstall CRS.
3. Execute the script `CRS Home/install/rootdeinstall.sh` on a local node to remove the OCR.
4. Run the OUI and in the Welcome page, click Deinstall Products to display the list of installed products on which you can select the CRS home to deinstall.

Note: A node is a local node if you plan to run Step 3 and 4 on that node.

Creating RAC Databases with the Database Configuration Assistant

This chapter describes how to use the Database Configuration Assistant (DBCA) in standalone mode to create and delete Real Application Clusters (RAC) databases. The topics in this chapter include:

- [Using the Database Configuration Assistant in Real Application Clusters](#)
- [Benefits of Using the Database Configuration Assistant](#)
- [Real Application Clusters High Availability Services](#)
- [Creating the Database after Installation Using the Database Configuration Assistant](#)
- [Creating a Real Application Clusters Database with the DBCA](#)
- [Deleting a Real Application Clusters Database with the DBCA](#)

See Also: *Oracle Real Application Clusters Administrator's Guide* for procedures on using the DBCA to add and delete instances

Using the Database Configuration Assistant in Real Application Clusters

The primary functions of the DBCA processing include:

- Create the database and its instances
- Set up network configuration for database, instances and database services
- Invoke EMCP and set up `targets.xml` file for Enterprise Manager Grid Control of the database
- Start up database, its instances, services, and any other node applications

See Also:

- ["Creating the Database after Installation Using the Database Configuration Assistant"](#) on page 9-3 for more information about using the DBCA in standalone mode
- The *Oracle Net Services Administrator's Guide* if you experience problems, for example, with the listener configuration, and for further information about Lightweight Directory Access Protocol (LDAP)-compliant directory support

Benefits of Using the Database Configuration Assistant

Oracle recommends that you use the DBCA to create your RAC database because the DBCA's preconfigured databases optimize your environment for features such as ASM, the server parameter file, and automatic undo management. The DBCA also provides pages to create new ASM disk groups if they are needed and configures recovery and backup disk space if you are using ASM or cluster file system storage.

With the DBCA, you may create site-specific tablespaces as part of database creation. If you have datafile requirements that differ from those offered by the DBCA templates, then create your database with the DBCA and modify the datafiles later. You can also execute user-specified scripts as part of your database creation process.

The DBCA also configures your RAC environment for various Oracle high availability features such as services and cluster administration tools. It also starts any database instances required to support your defined configuration.

Real Application Clusters High Availability Services

When you configure high availability services with the DBCA Database Services page, you can also configure service instance preferences and Transparent Application Failover (TAF) policies.

Service Configuration and Instance Preferences

Use the Database Services page button in the column labeled Not Used, Preferred, or Available to configure service instance preferences as described in the following list:

- Preferred—The service runs primarily on the selected instance
- Available—The service may run on the instance if a preferred instance fails
- Not Used—The service never runs on the instance

Transparent Application Failover Policies

Use the DBCA Database Services page to configure TAF failover policies. The DBCA Database Services page also has a TAF policy selector row under the instance preference display. Make a selection in this row for your failover and reconnection policy preference as described in the following list:

- None—Do not use TAF
- Basic—Establish connections at failover time
- Pre-connect—Establish one connection to a preferred instance and another connection to a backup instance that you have selected to be available

Creating the Database after Installation Using the Database Configuration Assistant

To create a database with the DBCA in standalone mode without ASM or a cluster file system, you must have configured each raw device as described in [Appendix C](#). In addition, you must have run the Oracle Net Configuration Assistant to configure your Oracle Net `listener.ora` file.

If you select DBCA templates that use preconfigured datafiles and if you do not use ASM or a cluster file system, then during database creation the DBCA first verifies that you created the raw devices for each tablespace. If you have not configured the raw devices, then you must configure the raw devices and replace the default datafile names that the DBCA provides with raw device names on the DBCA Storage page to continue database creation.

To start the DBCA, enter the `dbca` command from the `$ORACLE_HOME/bin` directory on one of the nodes.

Creating a Real Application Clusters Database with the DBCA

The first page that the DBCA displays is the Welcome page for RAC. The DBCA displays this RAC-specific Welcome page only if the Oracle home from which it is invoked was cluster installed.

See Also: The DBCA online help for more information

If the DBCA does not display the Welcome page for RAC, then the DBCA was unable to detect whether the Oracle home is cluster installed. In this case, check that the OUI inventory is correctly located in `/var/opt/oracle/oraInst.loc` and that `oraInventory` is not corrupted. Also, perform clusterware diagnostics by executing the `olsnodes` command from the `bin` directory in CRS home. Otherwise, to create a RAC database:

1. Select Real Application Clusters database, click Next, and the DBCA displays the Operations page. The DBCA enables the Configure Database Options, Delete a database, Instance Management, and Services Management options only if there is at least one RAC database configured on your cluster that runs from the Oracle home.
2. Select Create a database and click Next, and the DBCA displays the Node Selection page.
3. The DBCA highlights the local node by default. Select the other nodes that you want to configure as members of your cluster database, click Next, and the DBCA displays the Database Templates page. If nodes that are part of your cluster

installation do not appear on the Node Selection page, then perform inventory diagnostics and the clusterware diagnostics by executing the `olsnodes` command.

4. The templates on the Database Templates page are Custom Database, Transaction Processing, Data Warehouse, and General Purpose. The Custom Database template does not include datafiles or options specially configured for a particular type of application. Use one of the other templates, which include datafiles, if you want to create database with specifically configured options. Select the template from which you wish to create your cluster database, click Next, and the DBCA displays the Database Identification page.
5. Enter the global database name and the Oracle system identifier (SID) prefix for your cluster database, click Next, and the DBCA displays the Management Options page.

Note: The global database name can be up to 30 characters in length and must begin with an alphabetical character. The SID prefix must be between one and five characters in length and also begin with an alphabetical character. The DBCA uses the SID prefix to generate a unique value for the `ORACLE_SID` for each instance.

6. On the Management Options page, you can choose to manage your database with Enterprise Manager and choose either the Grid Control or Database Control option. If you select Enterprise Manager with the Grid Control option and DBCA discovers agents running on the local node, then you can select the preferred agent from a list. If you select the Database Control option, you can set up e-mail notification (for which you provide the outgoing mail server and e-mail address) and enable daily backups (for which you enter the backup time and OS credentials for the user that takes backup). After you make your selections and enter any required information, click Next, and the DBCA displays the Database Credentials page.
7. Enter the passwords for your database on the Database Credentials page. You can enter the same or different passwords for the users `SYS` and `SYSTEM`, plus `DBSNMP` and `SYSMAN` if you selected Enterprise Manager on the Management Options page. Select the Use the Same Password for All Accounts option to assign the same password to the listed users. Or provide a different password for each of these users by selecting the Use different Passwords option. Enter the password information, click Next, and DBCA displays the Storage Options page.
8. Use the Storage Options page to select a storage type for database creation. The Cluster File System option is the default. Select a storage option and click Next to proceed to the next page. If you selected Cluster File System, then the next page that appears is the Database File Locations page, described in Step 9. If you select Raw Devices, the next page that appears is the Recovery Configuration page, described in Step 10. If you select Automatic Storage Management (ASM), then provide additional information as follows:
 - a. If there is not an ASM instance on any of the cluster nodes, then the DBCA displays the Create ASM Instance page for you, described in Step c. Otherwise the DBCA displays the ASM Disk Groups page, described in Step d.
 - b. If an ASM instance exists on the local node, then the DBCA displays a dialog asking you to enter the password for the `SYS` user for ASM.

- c. To initiate the creation of the required ASM instance, supply the password for the SYS user of the ASM instance. If your Oracle home is installed on cluster file system, then the ASM instance uses an SPFILE, otherwise, you can select either an IFILE or an SPFILE on shared storage for the instances. After you enter the required information, click Next to create the ASM instance. Once the instance is created, DBCA proceeds to the ASM Disk Groups page described in Step d.
- d. From the ASM Disk Groups page, you can create a new disk group, add disks to an existing disk group, or select a disk group for your database storage.

If you have just created a new ASM instance, then there will be no disk groups from which to select, so you must create a new one by clicking Create New to open the Create Disk Group page, described in Step e.

Similarly, if one or more disk groups are displayed but you want to add a new one, then click Create New and follow the procedure described in Step e to complete the Create Disk Group page.

If you want to use an existing disk group but wish to add more disks to it, click Add Disks and follow the instructions in Step f.

Once you are satisfied with the ASM disk groups available to you, select the one you wish to use for your database files and click Next to proceed to the Database File Locations page, described in step 9.

Note: To use a flash recovery area, Oracle recommends that you create two separate ASM disk groups: one for the database files and one for the recovery area.

See Also: *Oracle Database Concepts* for more information about using a flash recovery area

- e. Enter the disk group name and then click the redundancy level for the group if you do not wish to use the default value (Normal). Create your disk group by selecting from the list of candidate disks. Continue by following the process described in Step g.
- f. If there is a disk group that you want to use but you want to add more disks to it, then select the group and click Add Disks. Add to the disk group by selecting from the list of candidate disks. Continue by following the process described in Step g.
- g. If you do not see the disks you wish to add, click Change Disk Discovery Path to alter the search path used by DBCA to find available disks. You can select disks with a status of Candidate or Former (never used in an ASM disk group or no longer in a group) just by checking the select box. To select a disk with a Member status, you must also check the Force column that is displayed to move the disk from its current ASM disk group to chosen disk group. When you have selected the desired disks, click OK to add them to your disk group and return to the ASM Disk Groups Page. To proceed, see the instructions described earlier in Step d.
- h. If the DBCA displays the following message:

The file `oracle_home/bin/oracle` does not exist on nodes `node_list`.
Make sure that file exists on these nodes before proceeding.

then the Oracle home from which the first ASM instance in the cluster runs is not installed on these cluster nodes. You must extend the ASM Oracle home to these nodes by performing the procedure documented in "Step 4: Adding Nodes at the Oracle RAC Database Layer" in the *Oracle Real Application Clusters Administrator's Guide*. However, do not perform Step 5 in that section. The OUI extends the ASM Oracle home to the selected nodes and performs any configuration required for running an ASM instance on these nodes.

- i. If the DBCA displays the message `Please run the DBCA from one of the nodes that has an existing ASM instance node_list.`, then you are attempting to create a RAC database using ASM storage and the ASM instance does not exist on the node from which you ran the DBCA. However, ASM instances do exist on the remote nodes that appear in the message node list. In this case, the DBCA cannot clone the existing ASM instance from the remote node to the local node. To correct this, start the DBCA from one of the nodes shown in the node list to create your RAC database using ASM storage. This copies the local node's ASM instance and modifies its parameters and attributes to create ASM instances on the nodes in your cluster that do not have ASM instances.
 - j. If the DBCA displays the message `ORACLE_BASE for ASM home asm_home does not match the value for database home db_home.` Please set `ORACLE_BASE` to `asm_home` and restart `DBCA.`, then this means that you selected a node to be part of your RAC database that does not have an ASM instance. In addition, the ASM instance on the local node is running from an Oracle home that is different from the Oracle home for the database to be created. Both the ASM home and the database home must be under a common `ORACLE_BASE`. If you created the original ASM instance without setting `ORACLE_BASE`, then set the `ORACLE_BASE` to the `asm_home` and restart the DBCA to proceed.
9. The Database File Locations page enables you to select the file storage for your database files: locations provided in a template, a common location for all database files (the files will not be Oracle-managed files), or Oracle-managed files in a common location. If you do not select the template option, you can enter an existing ASM disk group name or directory path name in the space provided, or click Browse to open a selection list.

If you wish to multiplex the database redo log files and control files, click `Multiplex Redo Logs and Control Files` and provide the location for each copy you want. Click OK when you have defined the multiplex locations to return to the Database File Locations page.

You may also define your own variables for the file locations if you plan to use the Database Storage page, explained in Step 14, to define individual file locations.

10. On the Recovery Configuration page, you can select redo log archiving by selecting `Enable Archiving`. If you are using ASM or CFS storage, then you can also select the flash recovery area and size on the Recovery Configuration page. If you are using ASM, then the flash recovery area defaults to the ASM Disk Group. If you are using CFS, then the flash recovery area defaults to `$ORACLE_BASE/flash_recovery_area`.

You may also define your own variables for the file locations if you plan to use the Database Storage page, explained in Step 14, to define individual file locations.

When you have completed your entries, click `Next`, and the Database Content page is displayed.

11. On the Database Content page, if you chose the Custom Database option, you can select or deselect the database components and their corresponding tablespace assignment. For a seed database, you can select whether to include the sample schemas in your database and to run custom scripts as part of the database creation processing. After completing your selections, click Next, and the Databases Services page is displayed.
12. To create services on the Database Services page, expand the Services tree. Oracle displays the global database name in the top left-hand corner of the page. Select the global database name and click Add to add services to the database. Enter a service name in the Add a Service dialog and click OK to add the service and return to the Database Services page.

The service name appears under the global database name. Select the service name and the DBCA displays the service preferences for the service on the right-hand side of the DBCA Database Services page. Change the instance preference (Not Used, Preferred, or Available) and TAF policies for the service as needed.

Repeat this procedure for each service and when you are done configuring services for your database, click Next. The DBCA displays the Initialization Parameters page.

13. By default, the Initialization Parameters page shows only the basic parameters and only enables you to change the parameter file definition if you are using raw storage. Each tab on the Initialization Parameters page provides different sets of information that you can add or modify as follows:

- a. **Memory Tab:** Click Typical for default values based on the database type you selected or Custom to set your own values for the memory parameters. You can also see values for the advanced parameters by clicking All Initialization Parameters.

Carefully review the parameter settings displayed in this dialog because the DBCA configures these settings in your server parameter file. Instance-specific parameter settings for your RAC database appear at the bottom of this dialog. The `sid` prefixes for these entries appear in the left column.

To review the instance-specific parameter settings, scroll downward using the scroll bar on the right side of the dialog. Use the check box in the Override Default column to indicate whether the DBCA should place the parameter setting in your server parameter file. The DBCA only places a parameter entry into the server parameter file if the entry displays a check mark in the Override Default column of the All Initialization Parameters dialog.

Notes:

- You cannot modify the `sid` in the Instance column
 - You can alter self-tuning parameters with this dialog. However, setting these parameters to inappropriate values may disable the Oracle self-tuning features.
 - You cannot specify instance-specific values for global parameters with the DBCA.
 - You should set the value of the `CLUSTER_DB_INSTANCES` parameter to the number of instances you intend to use in the cluster if you are not including all the related nodes during the current execution of DBCA.
 - If your global database name is longer than eight characters, then the database name value (in the `db_name` parameter) is truncated to the first eight characters and the `DB_UNIQUE_NAME` parameter value is set to the global name.
-
-

- b. Sizing Tab: Use this page to select the database standard block size and process count.
- c. Character Sets Tab: Use this page to set the database character set value.
- d. Connection Mode Tab: You can use this tab to select either dedicated or shared database connections to your database.
- e. Parameter File Tab: This tab will only appear if you are using raw storage. Use this tab to enter a raw device name for the location of the server parameter file.

When you have completed all your work on the Initialization Parameters page, click Next, and the Database Storage page is displayed.

14. If you selected a preconfigured database template, such as the General Purpose template, then the DBCA displays the control files, datafiles, and redo logs on the Database Storage page. Select the folder and the file name underneath the folder to edit the file name. However, if you selected the Custom Database template, the template without datafiles, then the DBCA displays the control files, tablespaces, datafiles, and redo logs. To change the tablespace properties, such as the datafile or the tablespace size, click the tablespaces icon to expand the object tree on the left-hand side of the page and click the tablespace. The tablespace property dialog appears on the right-hand side. Make your changes and click OK.

When entering file names in the Database Storage page for raw storage, if you have not set the `DBCA_RAW_CONFIG` environment variable, then the DBCA displays default datafile names. You must override these names to provide raw device names on this page for each control file, datafile, and redo log group file.

After you complete your entries on the Database Storage page, click Next, and the DBCA displays the Creation Options page.

15. On the Creation Options page, select one of the following database options and click Finish.
 - Create Database—Creates the database
 - Save as a Database Template—Creates a template that records the database structure, including user-supplied inputs, initialization parameters, and so on. You can later use this template to create a database.

- Generate Database Creation Scripts—Generates database creation scripts. The DBCA only displays this option if you selected the New Database template.

After you click Finish, the DBCA displays a Summary dialog.

16. Review the Summary dialog information and click OK to create the database.

After you complete Step 16 the DBCA:

- Creates an operative RAC database and its instances
- Creates the RAC data dictionary views
- Configures the network for the cluster database
- Starts the listeners and database instances and then starts the high availability services

Deleting a Real Application Clusters Database with the DBCA

This section explains how to delete a RAC database with the DBCA. This process deletes a database and removes a database's initialization parameter files, instances, OFA structure, and Oracle network configuration. However, this process does not remove datafiles if you placed the files on raw devices or on raw partitions.

To delete a database with the DBCA:

1. Start the DBCA on one of the nodes by executing the `dbca` command from the `$ORACLE_HOME/bin` directory. The DBCA Welcome page appears.
2. Select Oracle Real Application Clusters and click Next.

After you click Next, the DBCA displays the Operations page.

3. Select Delete a database, click Next, and the DBCA displays the List of Cluster Databases page.
4. If your user ID and password are not operating-system authenticated, then the List of Cluster Databases page displays the user name and password fields. If these fields appear, then enter a user ID and password that has SYSDBA privileges.

See Also: "Database Password and Role Management in Real Application Clusters" on page B-19

5. Select the database to delete and click Finish.

After you click Finish, the DBCA displays a Database Deletion Summary dialog showing the database name and associated instances that the DBCA is going to delete. This Summary dialog also describes the node on which each instance resides.

6. Click OK to begin the deletion of the database and its associated files, services, and environment settings, or click Cancel to stop the operation.

When you click OK, the DBCA continues the operation and deletes all of the associated instances for this database. The DBCA also removes the parameter files, password files, and `oratab` entries.

At this point, you have accomplished the following:

- Deleted the selected database from the cluster
- Deleted high availability services that were assigned to the database

- Deleted the Oracle Net configuration for the database
- Deleted the OFA directory structure from the cluster
- Deleted the datafiles if the datafiles were not on raw devices

Real Application Clusters Post-Installation Procedures

This chapter describes how to complete the post-installation tasks after you have installed the Oracle Database 10g with Real Application Clusters (RAC) software. It contains the following sections:

- [Required Post-Installation Tasks](#)
- [Recommended Post-Installation Tasks](#)

Note: This chapter only describes basic configurations. Refer to the *Oracle Database 10g Administrator's Guide for UNIX Systems* and the product administration and tuning guides for more detailed configuration and tuning information.

Required Post-Installation Tasks

You must perform the following tasks after completing your installation:

- [Download and Install Patches](#)
- [Configure Oracle Products](#)

Download and Install Patches

Refer to the *OracleMetaLink* Web site for required patches for your installation. To download required patches:

1. Use a Web browser to view the *OracleMetaLink* Web site:

`http://metalink.oracle.com`

2. Log in to *OracleMetaLink*.

Note: If you are not an *OracleMetaLink* registered user, then click Register for MetaLink! and register.

3. On the main *OracleMetaLink* page click Patches.
4. On the Select a Patch Search Area page click New MetaLink Patch Search.
5. On the Simple Search page click Advanced.
6. On the Advanced Search page click the search icon next to the Product or Product Family field.
7. In the Search and Select: Product Family window enter RDBMS Server in the For field and click Go.
8. Select RDBMS Server under the Results heading and click Select.
RDBMS Server appears in the Product or Product Family field and the current release appears in the Release field.
9. Select your platform from the list in the Platform field and click Go.
10. Any available patches appear under the Results heading.
11. Click the number of the patch that you want to download.
12. On the Patch Set page, click View README and read the page that appears. The README page contains information about the patch set and how to apply the patches to your installation.
13. Return to the Patch Set page, click Download, and save the file on your system.
14. Use the unzip utility provided with Oracle Database 10g to uncompress the Oracle patches that you downloaded from *OracleMetaLink*. The unzip utility is located in the `$ORACLE_HOME/bin` directory.

Configure Oracle Products

Many Oracle products and options must be configured before you use them for the first time. Before using individual Oracle Database 10g database products or options, refer to the manual in the product documentation library which is available on the documentation CD-ROM or on the OTN Web site.

Recommended Post-Installation Tasks

Oracle recommends that you perform the following tasks after completing an installation:

- [Back Up the root.sh Script](#)
- [Set Up Users Accounts](#)
- [Verifying Enterprise Manager Operations](#)

Back Up the root.sh Script

Oracle recommends that you back up the `root . sh` script after you complete an installation. If you install other products in the same Oracle home directory, then the Oracle Universal Installer (OUI) updates the contents of the existing `root . sh` script during the installation. If you require information contained in the original `root . sh` script, then you can recover it from the `root . sh` file copy.

Set Up Users Accounts

For information about setting up additional optional user accounts, refer to the *Oracle Database 10g Administrator's Guide for UNIX Systems*.

Verifying Enterprise Manager Operations

Run the following command to verify the Enterprise Manager configuration in your newly installed Real Application Clusters environment:

```
srvctl config database -d db_name
```

SRVCTL displays the name of the node and the instance for the node. The following example shows a node named `db1-server` running an instance named `db1` with a listener named `LISTENER_NODE1`. Execute the following command:

```
srvctl config database -d db
```

The output should be similar to:

```
db1-server db1 /private/system/db
db2-server db2 /private/system/db
```

When you complete these procedures, you are ready to perform the initial configuration tasks described in [Part IV](#).

Part IV

Real Application Clusters Environment Configuration

Part IV describes how to use the server parameter file (spfile) in Oracle Database 10g Real Application Clusters (RAC) and it describes the installed configuration. The chapters in Part IV are:

- [Chapter 11, "Configuring the Server Parameter File in Real Application Clusters Environments"](#)
- [Chapter 12, "Understanding the Real Application Clusters Installed Configuration"](#)

Configuring the Server Parameter File in Real Application Clusters Environments

This chapter describes server parameter file (spfile) placement and configuration in Real Application Clusters (RAC) environments. The topics in this chapter are:

- [Parameter Files and Real Application Clusters](#)
- [Using Server Parameter Files in Real Application Clusters](#)
- [Parameter File Search Order in Real Application Clusters](#)
- [Migrating to the Server Parameter File in Real Application Clusters Environments](#)
- [Server Parameter File Errors in Real Application Clusters](#)

See Also: The *Oracle Real Application Clusters Administrator's Guide* for more information about parameters and the *Oracle Real Application Clusters Deployment and Performance Guide* for a discussion of parallel execution-related parameters in RAC data warehouse environments

Parameter Files and Real Application Clusters

Oracle uses parameter settings in parameter files to determine how to control various database resources. You can use two types of files for parameter administration: the server parameter file (spfile) or one or more traditional client-side parameter files.

Oracle recommends that you administer parameters using spfile. If you use client-side parameter files, then Oracle does not preserve parameter changes made for self-tuning after shutdown.

See Also: The Oracle Database 10g Real Application Clusters documentation for more information about using client-side parameter files

Using Server Parameter Files in Real Application Clusters

By default, Oracle creates the server parameter file based on one `SPFILE`. You can change parameter settings in the server parameter file only by using Oracle Enterprise Manager or `ALTER SYSTEM SET SQL` statements; the server parameter file is a binary file that you should not edit.

Note: Oracle recommends that you avoid modifying the values for self-tuning parameters; overriding these settings can adversely affect performance.

If you are upgrading from a previous Oracle release, then create and configure the server parameter file for RAC using the procedures described in the following section.

Location of The Server Parameter File

The default location of the server parameter file when the database creates it from a `PFILE` is:

```
$ORACLE_HOME/dbs/spfile$ORACLE_SID.ora
```

The default location of the server parameter file is inappropriate for RAC databases if you use raw devices because all instances must use the same server parameter file.

Instead, Oracle recommends that you use a `PFILE` in the following directory:

```
$ORACLE_HOME/dbs/init$ORACLE_SID.ora
```

This path is valid for each instance and it refers to a single, shared initialization parameter file. The file must contain the following entry if you use raw storage:

```
SPFILE= '/dev/vx/rdisk/oracle_dg/dbspfile'
```

However, if you use a cluster file system, then use the following alternate file location:

```
SPFILE= '$ORACLE_HOME/dbs/spfile.ora'
```

If you use ASM, then the file location will be:

```
SPFILE= '+disk_group_name/dbunique_name/spfile<dbname>.ora'
```

You must use the same value of `SPFILE` so that all instances use the same server parameter file at startup.

To use the DBCA to create your database and to use the server parameter file, on the Initialization Parameters page select Create server parameter file (spfile) under the File Locations tab which is visible only if you are using raw storage. Then enter either a shared file system filename or the raw device path name in the `Server Parameters Filename` field.

Note: When you use the DBCA to create the server parameter file, the default PFILE file name is `$ORACLE_HOME/dbs/init$ORACLE_SID.ora`.

Parameter File Search Order in Real Application Clusters

Oracle searches for your parameter file in the following order:

1. `$ORACLE_HOME/dbs/spfilesid.ora`
2. `$ORACLE_HOME/dbs/spfile.ora`
3. `$ORACLE_HOME/dbs/initsid.ora`

Migrating to the Server Parameter File in Real Application Clusters Environments

Migrate to the server parameter file by creating and editing the server parameter file using the procedures described in this section.

Server Parameter File Placement in Real Application Clusters

For single-node cluster-enabled configurations, or if you are using a cluster file system, place the server parameter file on a file system. Otherwise, place the server parameter file on a shared raw device that is at least 5MB in size.

Procedures for Migrating to the Server Parameter File

Migrate to the server parameter file by completing the following procedures:

1. Combine the initialization parameter files for all instances into one `initdbname.ora` file by copying all shared IFILE contents *as is*. All parameters defined in your IFILE parameter files are global. Therefore, create them as "parameter=value" without `sid` prefixes.
2. Copy all instance-specific parameter definitions from `initsid.ora` files using the following syntax where `sid` is the `sid` of the instance:


```
sid.parameter=value
```
3. If you are using a cluster file system, then create the server parameter file using the `CREATE SPFILE` statement as in the following example:

```
CREATE SPFILE='?/dbs/spfile_dbname.ora'
FROM PFILE='?/dbs/initdbname.ora'
```

If you use ASM, then use the following syntax to create the server parameter file:

```
CREATE SPFILE='+disk_group_name/db_uniquename/spfiledbname.ora'
FROM PFILE='?/dbs/initdbname.ora'
```

If you use raw storage, then use the following syntax to create the server parameter file on a raw device:

```
CREATE SPFILE='/dev/vx/rdisk/oracle_dg/dbspfile'  
FROM PFILE='?/db/initdbname.ora'
```

These statements read your combined `initdbname.ora` file that you created by merging your `IFILES`. Then it transfers the settings for the parameters from the merged file into your server parameter file.

4. Oracle recommends that you use the server parameter file by executing the `STARTUP` command as in this example:

```
STARTUP PFILE=$ORACLE_HOME/db/initsid.ora
```

Where the file `initsid.ora` contains the entry:

```
SPFILE='/dev/vx/rdisk/oracle_dg/dbspfile'
```

If you use this `STARTUP` command syntax, then Oracle uses the server parameter file entry specified in `initsid.ora`.

Server Parameter File Errors in Real Application Clusters

Oracle reports errors that occur during server parameter file creation or while reading the file during startup. If an error occurs during a parameter update, then Oracle records the error in your `ALERT.LOG` file and ignores subsequent parameter updates to the file. If this happens, then do either of the following:

- Shut down the instance, recover the server parameter file, and restart the instance.
- Enable the instance to continue running without regard for subsequent parameter updates.

Oracle displays errors for parameter changes that you attempt when you incorrectly use the `ALTER SYSTEM SET` statement. Oracle does this when an error occurs while reading from or writing to the server parameter file.

See Also: *The Oracle Real Application Clusters Administrator's Guide* for more information about backing up the `SPFILE`

Understanding the Real Application Clusters Installed Configuration

This chapter describes the Real Application Clusters (RAC) installed configuration. The topics in this chapter include:

- Understanding the Configured Environment in Real Application Clusters
- The Oracle Cluster Registry in Real Application Clusters
- oratab Configurations for Real Application Clusters
- Database Components Created Using the Database Configuration Assistant
- Managing Undo Tablespaces in Real Application Clusters
- Initialization Parameter Files
- Configuring Service Registration-Related Parameters in Real Application Clusters
- Configuring the Listener File (listener.ora)
- Directory Server Access (ldap.ora File)
- Net Service Names (tnsnames.ora File)
- Profile (sqlnet.ora File)

Understanding the Configured Environment in Real Application Clusters

The Oracle Net Configuration Assistant (NetCA) and the Database Configuration Assistant (DBCA) configure your environment to meet the requirements for database creation and Enterprise Manager discovery of Real Application Cluster databases.

Note: Configuration files are created on each node in your cluster database.

The Oracle Cluster Registry in Real Application Clusters

The Database Configuration Assistant uses the Oracle Cluster Registry (OCR) for storing the configurations for the cluster databases that it creates. The OCR is a shared file in a cluster file system environment. If you do not use a cluster file system, then you must make this file a shared raw device. The `CRS_home/root.sh` script initializes the OCR during the CRS installation.

oratab Configurations for Real Application Clusters

Oracle creates an entry for each RAC database in the `oratab` configuration file. Oracle Enterprise Manager uses this file during service discovery to determine the name of the RAC database as well whether the database should be automatically started upon restart. The database entry has the following syntax:

```
db_unique_name:$ORACLE_HOME:N
```

where `db_unique_name` is the database name for your RAC database, `$ORACLE_HOME` is the directory path to the database, and `N` indicates that the database should not be started at restart time. A sample entry for a database named `db` is:

```
db:/private/system/db:N
```

Note: Where the notation `db_name` appears in the previous example and throughout this chapter, it refers to the database name you entered when prompted by the DBCA, or it refers to the entry you made for the `DATABASE` keyword of the `CREATE DATABASE` statement.

Database Components Created Using the Database Configuration Assistant

This section describes the database components that the DBCA creates which include:

- [Tablespaces and Datafiles](#)
- [Control Files](#)
- [Redo Log Files](#)

Tablespaces and Datafiles

An Oracle database for both single-instance and cluster database environments is divided into smaller logical areas of space known as tablespaces. Each tablespace

corresponds to one or more datafiles stored on a disk. [Table 12–1](#) shows the tablespace names used by a RAC database and the types of data they contain:

Table 12–1 Tablespace Names that Real Application Clusters Databases Use

Tablespace Name	Contents
SYSTEM	Consists of the data dictionary, including definitions of tables, views, and stored procedures needed by the database. Oracle automatically maintains information in this tablespace.
SYSAUX	An auxiliary system tablespace that contains the DRSYS (contains data for OracleText), CWMLITE (contains the OLAP schemas), XDB (for XML features), ODM (for Oracle Data Mining), TOOLS (contains Enterprise Manager tables), INDEX, EXAMPLE, and OEM-REPO tablespaces.
USERS	Consists of application data. As you create and enter data into tables, Oracle fills this space with your data.
TEMP	Contains temporary tables and indexes created during SQL statement processing. You may need to expand this tablespace if you are executing a SQL statement that involves significant sorting, such as ANALYZE COMPUTE STATISTICS on a very large table, or the constructs GROUP BY, ORDER BY, or DISTINCT.
UNDOTBS <i>n</i>	These are the undo tablespaces for each instance that the DBCA creates for automatic undo management.
RBS	If you do not use automatic undo management, then Oracle uses the RBS tablespace for the rollback segments.

You cannot alter these tablespace names when using the preconfigured database configuration options from the Oracle Universal Installer. However, you can change the names of the tablespaces if you use the customized database creation method.

As mentioned, each tablespace has one or more datafiles and you are prompted for the names of the datafiles for a preconfigured database.

Control Files

The database is configured with two control files that are stored on shared storage.

Redo Log Files

Each instance is configured with at least two redo log files that are stored in the shared storage. If you chose cluster file system, then these files are shared file system files. If you do not have a cluster file system, then these files are raw devices. If you use ASM, then these files are stored on the ASM disk group.

The file names of the redo log files that are created with the preconfigured database configuration options vary by storage type. You must enter the raw device names unless you are using a cluster file system.

When using raw devices, to use the customized database creation method, locate the redo log files in the Database Storage page and replace their default filenames with the correct raw device names or symbolic link names.

Managing Undo Tablespaces in Real Application Clusters

Oracle stores rollback or undo information in undo tablespaces. To manage undo tablespaces, Oracle recommends that you use automatic undo management. Automatic undo management is an automated undo tablespace management mode that is easier to administer than manual undo management.

See Also: *Oracle Real Application Clusters Administrator's Guide* for more information about managing undo tablespaces

Initialization Parameter Files

Oracle recommends using the server parameter file (spfile). This file resides on the server on the shared disk; all instances in a cluster database can access this parameter file.

See Also: [Chapter 11, "Configuring the Server Parameter File in Real Application Clusters Environments"](#) for more information about the creation and use of parameter files

Configuring Service Registration-Related Parameters in Real Application Clusters

Two key benefits of RAC are connection load balancing and failover. RAC extends the ability of single-instance Oracle database load balancing, where connections are distributed among local dispatchers, to the balancing of connections among all instances in a cluster database. In addition, RAC provides failover by configuring multiple listeners on multiple nodes to manage client connection requests for the same database service. Connection load balancing and failover increase availability by taking advantage of the redundant resources within a cluster database. These features, however, require cross instance registration.

Cross instance registration in RAC occurs when an instance's PMON process registers with the local listener and with all other listeners. Thus, all instances in the cluster database register with all listeners that run on nodes that run instances of the cluster database. This enables all listeners to manage connections across all instances for both load balancing and failover.

Cross instance registration requires configuring the `LOCAL_LISTENER` and `REMOTE_LISTENER` initialization parameters. The `LOCAL_LISTENER` parameter identifies the local listener and the `REMOTE_LISTENER` parameter identifies the global list of listeners. The `REMOTE_LISTENER` parameter is dynamic. Oracle changes the setting for `REMOTE_LISTENER` dynamically when you reconfigure your cluster database, for example, when you add or delete instances.

By default, the DBCA configures your environment with dedicated servers. However, if you select the Shared server option on the DBCA, then Oracle configures the shared server. In this case, Oracle uses both dedicated and shared server processing. When shared servers are configured, the `DISPATCHERS` parameter is specified as in the following example:

```
DISPATCHERS="(protocol=tcp)"
```

If the `DISPATCHERS` initialization parameter does not specify the `LISTENER` attribute as in the previous example, then the PMON process registers information for all dispatchers with the listeners specified by the `LOCAL_LISTENER` and `REMOTE_LISTENER` parameters.

However, when the `LISTENER` attribute is specified, the PMON process registers dispatcher information with the listeners specified by the `LISTENER` attribute. In this case, setting the `LISTENER` attribute overrides `REMOTE_LISTENER` settings for the specified dispatchers as in the following example:

```
DISPATCHERS="(protocol=tcp)(listener=listeners_db_name)"
```

See Also: *Oracle Net Services Administrator's Guide* for further information about cross instance registration, shared and dedicated server configurations, and connection load balancing

Configuring the Listener File (listener.ora)

You can configure two types of listeners in the `listener.ora` file as described under the following headings:

- [Local Listeners](#)
- [Multiple Listeners](#)
- [How Oracle Uses the Listener \(listener.ora File\)](#)

Local Listeners

If you configured dedicated server mode using the DBCA Connection Mode tab on the Initialization Parameters page, then DBCA automatically configures the `LOCAL_LISTENER` parameter when the listener uses a non-default address port.

If you configured the dedicated server mode by setting the `REMOTE_LISTENER` initialization parameter, then you must also configure the instance-specific `LOCAL_LISTENER` initialization parameter.

For example, to configure the `LOCAL_LISTENER` parameter, add the following entry to the initialization parameter file, where `listener_sid` is resolved to a listener address through either a `tnsnames.ora` file or through the Oracle Names Server:

```
sid.local_listener=listener_sid
```

The following entry should be in your `tnsnames.ora` file:

```
listener_sid=(address=(protocol=tcp)(host=node1-vip)(port=1522))
```

Multiple Listeners

If the DBCA detects more than one listener on the node, it displays a list of the listeners. You can select one or all of these listeners with which to register your database.

How Oracle Uses the Listener (listener.ora File)

Services coordinate their sessions using listener file entries by running a process on the server that receives connection requests on behalf of a client application. Listeners are configured to respond to connection requests sent to protocol addresses for a database service or non-database service.

Protocol addresses are configured in the listener configuration file, `listener.ora`, for a database service or a non-database service. Clients configured with the same addresses can connect to a service through the listener.

During a preconfigured database configuration installation, the Oracle Net Configuration Assistant creates and starts a default listener called `LISTENER_NODENAME`. The listener is configured with a default protocol listening addresses for the database and external procedures. The customized installation process prompts you to create at least one listener with the Oracle Net Configuration Assistant. The listener is configured to respond to connection requests that are directed at one protocol address you specify, as well as an address for external procedures.

Both installation modes configure service information about the RAC database and external procedures. An Oracle Database 10g Release 1 (10.1) database service automatically registers its information with the listener, such as its service name, instance names, and load information.

This feature, called service registration, does not require configuration in the `listener.ora` file. After listener creation, the Oracle Net Configuration Assistant starts the listener. A sample `listener.ora` file with an entry for an instance named `node1` is:

```
listener_node1=
  (description=
    (address=(protocol=ipc)(key=extproc))
    (address=(protocol=tcp)(host=node1-vip)(port=1521))
    (address=(protocol=tcp)(host=node1-ip)(port=1521)))
sid_list_listener_node1=
  (sid_list=
    (sid_desc=
      (sid_name=plsextproc)
      (oracle_home=/private/system/db)
      (program=extproc)
```

Listener Registration and PMON Discovery

When a listener starts after the Oracle instance starts, and the listener is listed for service registration, registration does not occur until the next time the PMON discovery routine executes. By default, PMON discovery occurs every 60 seconds.

To override the 60-second delay, use the SQL `ALTER SYSTEM REGISTER` statement. This statement forces PMON to register the service immediately.

Oracle recommends that you create a script to execute this statement immediately after starting the listener. If you execute this statement while the listener is up and the instance is already registered, or while the listener is down, then the statement has no effect.

See Also: *Oracle Net Services Administrator's Guide* for further information about the listener and the `listener.ora` file

Directory Server Access (ldap.ora File)

If you configure access to a Lightweight Directory Access Protocol (LDAP)-compliant directory server with the Oracle Net Configuration Assistant during a customized installation, an `ldap.ora` file is created. The `ldap.ora` file contains the following types of information:

- Type of directory
- Location of the directory
- Administrative context from which this server can look up, create, and modify a net service name, and the database service entries

See Also: *Oracle Net Services Administrator's Guide* for further information about directory naming configuration and directory server access configuration

Net Service Names (tnsnames.ora File)

A `tnsnames.ora` file is created on each node with net service names. A connect identifier is an identifier that maps to a connect descriptor. A connect descriptor contains the following information:

- The network route to the service, including the location of the listener through a protocol address
- The `SERVICE_NAME` for an Oracle release 8.1 or later, or `SID` for pre-8.1 Oracle releases

Note: The `SERVICE_NAME` parameter you use in `tnsnames.ora` is singular because you can only specify one service name.

The DBCA creates net service names for connections as shown in [Table 12-2](#):

Table 12-2 Connections for Net Service Names

Net Service Name Type	Description
Database connections	<p>Clients that connect to any instance of the database use the net service name entry for the database. This entry also enables Oracle Enterprise Manager to discover a RAC database.</p> <p>A listener address is configured for each node that runs an instance of the database. The <code>LOAD_BALANCE</code> option causes Oracle to choose the address randomly. If the chosen address fails, then the <code>FAILOVER</code> option causes the connection request to fail over to the next address. Thus, if an instance fails, then clients can still connect using another instance.</p> <p>In the following example, <code>db.us.oracle.com</code> is used by the client to connect to the target database, <code>db.us.oracle.com</code>.</p> <pre>db.us.acme.com= (description= (load_balance=on) (address=(protocol=tcp)(host=node1-vip)(port=1521) (address=(protocol=tcp)(host=node2-vip)(port=1521) (connect_data= (service_name=db.us.acme.com)))</pre> <p>Note: <code>FAILOVER=ON</code> is set by default for a list of addresses. Thus, you do not need to explicitly specify the <code>FAILOVER=ON</code> parameter.</p> <p>When you set <code>DB_UNIQUE_NAME</code> by entering a global database name that is longer than eight characters, excluding <code>DB_DOMAIN</code>, then a net service entry similar to the following is created:</p> <pre>mydatabase.us.acme.com= (description = (address = (protocol = tcp)(host = node1-vip)(port = 1521)) (address = (protocol = tcp)(host = node2-vip)(port = 1521)) (load_balance = yes) (connect_data = (server = dedicated) (service_name = mydatabase.us.acme.com)))</pre>

Table 12–2 (Cont.) Connections for Net Service Names

Net Service Name Type	Description
Instance connections	<p>Clients that connect to a particular instance of the database use the net service name entry for the instance. This entry, for example, enables Oracle Enterprise Manager to discover the instances in the cluster. These entries are also used to start and stop instances.</p> <p>In the following example, <code>db1.us.acme.com</code>, is used by Oracle Enterprise Manager to connect to an instance named <code>db1</code> on <code>db1-server</code>:</p> <pre>db1.us.acme.com= (description= (address=(protocol=tcp)(host=node1-vip)(port=1521)) (connect_data= (service_name=db.us.acme.com) (instance_name=db1)))</pre>
Remote listeners	<p>As discussed in "Configuring Service Registration-Related Parameters in Real Application Clusters" on page 12-4, the <code>REMOTE_LISTENER</code> parameter identifies the global list of listeners and it is dynamic. Oracle changes the setting for <code>REMOTE_LISTENER</code> when you reconfigure your cluster database.</p> <p>Whether using shared servers or dedicated servers, the list of remote listeners is supplied using the <code>REMOTE_LISTENERS</code> parameter, for example:</p> <pre>REMOTE_LISTENERS=listeners_db_unique_name</pre> <p>This enables the instance to register with remote listeners on the other nodes; <code>listeners_db_unique_name</code> is resolved through a naming method such as a <code>tnsnames.ora</code> file.</p> <p>In the following example, <code>listeners_db.us.acme.com</code> is resolved to a list of listeners available on the nodes on which the cluster database has instances:</p> <pre>listeners_db.us.acme.com= (address_list= (address=(protocol=tcp)(host=node1-vip)(port=1521)) (address=(protocol=tcp)(host=node2-vip)(port=1521)))</pre> <p>The instance uses this list to determine the addresses of the remote listeners with which to register its information.</p>
Nondefault listeners	<p>As discussed in "Local Listeners" on page 12-5 and "Multiple Listeners" on page 12-5, the <code>LOCAL_LISTENER</code> parameter is set in the <code>init_{sid}.ora</code> file if a non-default listener is configured, for example:</p> <pre>sid.local_listener=listener_sid</pre> <p>Where <code>listener_{sid}</code> is resolved to a listener address through a naming method such as a <code>tnsnames.ora</code> file.</p> <p>In the following sample, <code>listener_db1.us.acme.com</code> is resolved to the nondefault listener address:</p> <pre>listener_db1.us.acme.com= (address=(protocol=tcp)(host=node1-vip)(port=1522))</pre>

Table 12–2 (Cont.) Connections for Net Service Names

Net Service Name	Type	Description
Services Entries		<p>When you configure high availability services using the DBCA Services page, then the DBCA creates net service entries similar to the following. The three services in the following examples, db_svc1, db_svc2, and db_svc3, have TAF policies of NONE, BASIC and PRECONNECT respectively.</p> <pre> db_svc1.us.acme.com= (description = (address=(protocol=tcp)(host=node1-vip)(port=1521)) (address=(protocol=tcp)(host=node2-vip)(port=1521)) (load_balance=yes) (connect_data= (server = dedicated) (service_name = db_svc1.us.acme.com))) db_svc2.us.acme.com= (description= (address=(protocol=tcp)(host=node1-vip)(port=1521)) (address=(protocol=tcp)(host=node2-vip)(port=1521)) (load_balance=yes) (connect_data = (server = dedicated) (service_name=db_svc2.us.acme.com) (failover_mode = (type=select) (method=basic) (retries=180) (delay=5)))) db_svc3.us.acme.com= (description= (address=(protocol=tcp)(host=node1-vip)(port=1521)) (address=(protocol=tcp)(host=node2-vip)(port=1521)) (load_balance=yes) (connect_data= (server=dedicated) (service_name=db_svc3.us.acme.com) (failover_mode= (backup=db_svc3_preconnect.us.acme.com) (type=select) (method=preconnect) (retries=180) (delay=5)))) </pre>

Table 12–2 (Cont.) Connections for Net Service Names

Net Service Name Type	Description
Services Entries (continued)	<p>When a service has a TAF policy of PRECONNECT, then a <code>service_name_preconnect</code> net service entry is also created as in the following example:</p> <pre>db_svc3_preconnect.us.acme.com = (description = (address = (protocol = tcp)(host = node1-vip)(port = 1521)) (address = (protocol = tcp)(host = node2-vip)(port = 1521)) (load_balance = yes) (connect_data = (server = dedicated) (service_name = ktest_svc3_preconnect.us.amce.com) (failover_mode = (backup = db_svc3.us.acme.com) (type = select) (method = basic) (retries = 180) (delay = 5)))))</pre>
External procedures	<p>An entry for connections to external procedures. This enables an Oracle Database 10g database to connect to external procedures.</p> <pre>extproc_connection_data.us.acme.com= (description= (address_list= (address=(protocol=ipc)(key=extproc0))) (connect_data= (sid=plsextproc)))</pre>

Example 12–1 Example tnsnames.ora File

The following is a sample `tnsnames.ora` file that is created during a preconfigured database configuration installation:

```
db.us.acme.com=
  (description=
    (load_balance=on)
    (address=(protocol=tcp)(host=node1-vip)(port=1521))
    (address=(protocol=tcp)(host=node2-vip)(port=1521))
    (connect_data=
      (service_name=db.us.acme.com)))

db1.us.acme.com=
  (description=
    (address=(protocol=tcp)(host=node1-vip)(port=1521))
    (connect_data=
      (service_name=db.us.acme.com)
      (instance_name=db1)))

db2.us.acme.com=
  (description=
    (address=(protocol=tcp)(host=node2-vip)(port=1521))
    (connect_data=
      (service_name=db.us.acme.com)
      (instance_name=db2)))

listeners_db.us.acme.com=
  (address_list=
```

```

(address=(protocol=tcp)(host=node1-vip)(port=1521))
(address=(protocol=tcp)(host=node2-vip)(port=1521))

extproc_connection_data.us.acme.com=
(description=
(address_list=
(address=(protocol=ipc)(key=extproc)))
(connect_data=
(sid=plsextproc)
(presentation=R0)))

```

See Also: *Oracle Net Services Administrator's Guide* for further information about the `tnsnames.ora` file

Profile (sqlnet.ora File)

The `sqlnet.ora` file is automatically configured with:

- The computer's domain
 - This domain is automatically appended to any unqualified net service name. For example, if the default domain is set to `us.acme.com`, then Oracle resolves `db` in the connect string `CONNECT scott/tiger@db as: db.us.acme.com`.
- A naming method the server uses to resolve a name to a connect descriptor
 - The order of naming methods is as follows: directory naming (for the customized installations only), `tnsnames.ora` file, Oracle Names server, and host naming.

The following is a sample `sqlnet.ora` file created during a preconfigured database configuration install:

```

names.default_domain=us.acme.com
names.directory_path=(tnsnames, onames,hostname)

```

See Also: *The Oracle Net Services Administrator's Guide* for further information about the `sqlnet.ora` file

Part V

Real Application Clusters Installation and Configuration Reference Information

Part V provides Real Application Clusters (RAC) installation and configuration reference information. The contents of Part V are:

- [Appendix A, "Troubleshooting the Real Application Clusters Installation Process"](#)
- [Appendix B, "Using Scripts to Create Real Application Clusters Databases"](#)
- [Appendix C, "Configuring Raw Devices for Real Application Clusters"](#)
- [Appendix D, "Converting to Real Application Clusters from Single-Instance Oracle Databases"](#)
- [Appendix E, "Directory Structure for Oracle Database 10g Real Application Clusters Environments"](#)

Troubleshooting the Real Application Clusters Installation Process

This appendix provides troubleshooting information for installing Oracle Database 10g Real Application Clusters (RAC). The topic in this appendix is:

- [Troubleshooting the Real Application Clusters Installation](#)

See Also: The Oracle Database 10g Real Application Clusters documentation set included on your Oracle Database 10g Server Documentation CD:

- *Oracle Real Application Clusters Administrator's Guide*
- *Oracle Real Application Clusters Deployment and Performance Guide*

Troubleshooting the Real Application Clusters Installation

This section contains these topics:

- [Real Application Clusters Installation Error Messages](#)
- [Performing Cluster Diagnostics During Real Application Clusters Installations](#)

Real Application Clusters Installation Error Messages

Real Application Clusters Management Tools Error Messages are in *Oracle Real Application Clusters Administrator's Guide*.

Performing Cluster Diagnostics During Real Application Clusters Installations

If the Oracle Universal Installer (OUI) does not display the Node Selection page, then perform clusterware diagnostics by executing the `olsnodes -v` command from the `CRS home/bin` directory and analyzing its output. Refer to your clusterware documentation if the detailed output indicates that your clusterware is not running.

Using Scripts to Create Real Application Clusters Databases

This chapter describes the steps required to create an Oracle Real Application Clusters (RAC) database from scripts. The topic in this appendix is:

- [Creating a Database Using Scripts](#)

Creating a Database Using Scripts

To generate a script to create a Real Application Clusters database, create a database from the script, and prepare the database for use, execute the following steps:

1. Invoke the Database Configuration Assistant (DBCA) and select your preferred options to build the RAC database. However, you must select the Custom Database template on the Database Templates page for the DBCA to provide the script generation option.

On the Creation Options page of your DBCA session, deselect Create Database and select Generate Database Creation Scripts before you click Finish. You can accept the default destination directory for the script file or browse for a different location. In either case, you should note the path name for use in the next step.

See Also: ["Creating RAC Databases with the Database Configuration Assistant"](#) for details on executing a DBCA session.

2. Navigate to the directory where the DBCA stored your script and review the SQL scripts that are there to ensure that they contain the statements to build a database with the characteristics you require. If they do not, Oracle recommends that you rerun the DBCA to create scripts with the desired configuration rather than editing the scripts yourself.
3. Execute the `db_name.sql` script, where `db_name` is the database name that you assigned in your earlier DBCA session, using a SQL*Plus session with the operating system SID set to the value you provided in your DBCA session. This script executes the `CreatedB.sql`, `postDBCcreation.sql`, `CreatedBCatalog.sql`, and `CreatedBFiles.sql` scripts followed by the scripts to configure the optional products that you selected.
4. Run SVRCTL to configure instance characteristics and node applications as described in the *Oracle Real Application Clusters Administrator's Guide*.
5. Configure Net Services to support your new database and instances as described in [Chapter 12, "Understanding the Real Application Clusters Installed Configuration"](#).

Configuring Raw Devices for Real Application Clusters

This appendix provides additional information about configuring raw devices to deploy Real Application Clusters (RAC). You must configure raw devices if you do not use ASM or a cluster file system. The topic in this appendix is:

- [Configuring Shared Disk Subsystems for Real Application Clusters](#)

Configuring Shared Disk Subsystems for Real Application Clusters

If you select one of the preconfigured database options on the Oracle Universal Installer's Database Configuration page, or if you use the Database Configuration Assistant (DBCA) interactively to create a preconfigured database and you do not use a cluster file system, then you must configure additional raw devices as described under the heading "[Additional Raw Devices Required by the DBCA for Non-CFS Environments](#)" on page C-2.

Additional Raw Devices Required by the DBCA for Non-CFS Environments

Configure raw devices as described in this section to use the DBCA to create a database on raw devices. These devices are in addition to the OCR and voting disk mentioned in the previous section. Create these devices before running the OUI to install the Oracle Database 10g software. The DBCA cannot create a RAC database unless you have properly configured the following devices:

- Four raw devices for four tablespace datafiles
- At least two raw devices for control files
- One raw device for each instance for its own tablespace for automatic undo management, or one rollback segment tablespace to use manual undo management
- At least two raw devices for redo log files for each instance
- One raw device for the server parameter file

Note: Each instance has its own redo log files, but all instances in a cluster share the control files and datafiles. In addition, each instance's log files must be readable by all other instances for recovery.

Planning Your Raw Device Creation Strategy

Before installing the Oracle Database 10g software with RAC, create enough partitions of specific sizes to support your database and also leave a few spare partitions of the same size for future expansion. For example, if you have space on your shared disk array, select a limited set of standard partition sizes for your entire database. Partition sizes of 50MB, 100MB, 500MB, and 1GB are suitable for most databases. Also create a few very small and a few very large spare partitions that are, for example, 1MB and perhaps 5GB or greater in size. Based on your plans for using each partition, determine the placement of these spare partitions by combining different sizes on one disk, or by segmenting each disk into same-sized partitions.

Note: Ensuring that there are spare partitions enables you to perform emergency file relocations or additions if a tablespace datafile becomes full.

Converting to Real Application Clusters from Single-Instance Oracle Databases

This chapter describes the procedures for converting from single-instance Oracle databases to Real Application Clusters (RAC) databases. The topics in this appendix are:

- [Deciding to Convert](#)
- [Prerequisites for Conversion](#)
- [Single-Instance to Cluster-Enabled Conversion Administrative Issues](#)
- [Converting from Single-Instance to Real Application Clusters](#)
- [Post-Conversion Steps](#)

If you are upgrading from Oracle Parallel Server to RAC or from an earlier version of RAC, then use the Database Upgrade Assistant (DBUA). In other words, the procedures in this chapter assume that your original single-instance database and the target RAC database are of the same version and running on the same platform.

See Also: *Oracle Database Upgrade Guide* for more information about the DBUA

Deciding to Convert

Do not convert to RAC if:

- You are not using a supported configuration of a cluster file system or shared disks
- Your application was specifically designed to *not* use cluster database processing

If your platform supports a cluster file system, then you can use it for RAC. You can also convert to RAC and use a non-shared file system. In either case, Oracle strongly recommends that you use Oracle Universal Installer (OUI) to perform an Oracle Database 10g installation that sets up the Oracle home in an identical location on each of the selected nodes in your cluster.

Prerequisites for Conversion

Your system must meet the following hardware and software requirements to convert to RAC:

- A supported hardware and operating system software configuration
- An additional license for Oracle Database 10g with RAC

Single-Instance to Cluster-Enabled Conversion Administrative Issues

Note the following administrative considerations before conversion:

- Backup procedures should be available before converting from a single-instance Oracle database to RAC.
- Additional archiving considerations apply in RAC environments. In particular, the archive file format requires a thread number. In addition, the archived logs from all instances of a RAC database are required for media recovery. If you archive to a file and you do not use a cluster file system, then a method of accessing the archive logs from all nodes on which the cluster database has instances is required where file systems are not shared.

Converting from Single-Instance to Real Application Clusters

To convert from single-instance Oracle databases to RAC, Oracle strongly recommends that you use the Database Configuration Assistant (DBCA). This is because the DBCA automates the configuration of the control file attributes, creates the undo tablespaces and the redo logs, and makes the initialization parameter file entries for cluster-enabled environments. It also configures the Oracle Net Services, Cluster Ready Services (CRS) resources, and the configuration for RAC database management for use by Oracle Enterprise Manager or the SRVCTL utility. This section describes the following scenarios:

- [Single Instance on a Non-Cluster Machine to Oracle Database 10g with RAC](#)
- [Single Instance on a Cluster to Oracle Database 10g RAC](#)

Single Instance on a Non-Cluster Machine to Oracle Database 10g with RAC

To convert from a single-instance Oracle database that is on a non-cluster machine to RAC, perform the procedures described under the following headings in the order shown:

- [Back up the Original Single-Instance Database](#)

- [Perform the Pre-Installation Steps](#)
- [Set up the Cluster](#)
- [Copy the Preconfigured Database Image](#)
- [Install Oracle Database 10g Software with Real Application Clusters](#)

Back up the Original Single-Instance Database

Use the DBCA to create a preconfigured image of your single-instance database by invoking the DBCA from the `bin` directory under `ORACLE_HOME` and choosing Welcome > Manage Templates > Create a database template [select From an existing database (structure as well as data)] > Database instance [select the database instance name] > Template Name [enter template name, use database instance as the default, and description] > Finish.

The DBCA will generate two files, a database structure file (`template_name.dbc`) and database preconfigured image file (`template_name.dfb`). These files are generated by default in the `ORACLE_HOME/assistants/dbca/templates` directory.

Perform the Pre-Installation Steps

Perform the pre-installation steps as documented in [Part II](#) of this book. This includes creating the `oracle` user account and the `dba` group on all nodes, setting up `oracle` user equivalence, set the `DBCA_RAW_CONFIG` environment variable, and so on. Then set up shared storage by referring to the "Configure Disk Storage for Oracle Database and Recovery Files" sections in the pre-installation chapters in [Part II](#).

See Also: Storage vendor-specific documentation for setting up the shared disk subsystem and for information about how to mirror and stripe disks

Set up the Cluster

To use vendor clusterware, form a cluster with the required number of nodes according to your vendor's documentation. Once you have configured all of the nodes in your cluster, either with or without vendor clusterware, install CRS by referring to the procedures in [Chapter 7, "Installing Cluster Ready Services"](#).

Copy the Preconfigured Database Image

This includes copying the database structure `*.dbc` file and the database preconfigured image `*.dfb` file that the DBCA created in the previous procedure "[Back up the Original Single-Instance Database](#)" on page D-3 to a temporary location on the node in the cluster from which you plan to run the DBCA.

Install Oracle Database 10g Software with Real Application Clusters

1. Run the Oracle Universal Installer (OUI) to perform an Oracle installation with the Oracle 10g Database with RAC.
2. Select Cluster Installation Mode on the Specify Hardware Cluster Installation page of the Oracle Universal Installer (OUI) and select the nodes to include in your RAC database.
3. On the OUI Database Configuration Types page, select the Customized install type and use the template that you copied to a temporary location in the "[Copy the Preconfigured Database Image](#)" procedure. Use the browse option to select the template location.

After installing the Oracle software, the OUI runs the post installation configuration tools such as the Network Configuration Assistant (NetCA), the DBCA, and so on.

4. If you selected raw storage on the OUI Storage Options page, then on the DBCA File Locations Tab on the Initialization Parameters page, replace the data files, control files, and log files, and so on, with the corresponding raw device files if you have not setup the `DBCA_RAW_CONFIG` environment variable. You must also make these entries on the Storage page.

See Also: [Chapter 9](#) for more details about the DBCA

5. After creating the RAC database, the DBCA displays the Password Management page on which you must change the passwords for database privileged users who have `SYSDBA` and `SYSOPER` roles. Once the DBCA exits, the conversion process is complete.

Single Instance on a Cluster to Oracle Database 10g RAC

There are three scenarios in which a single-instance database can exist on a cluster machine:

- Scenario 1: The Oracle home from which the single-instance database is running is cluster installed.
- Scenario 2: The Oracle home from which the single-instance database is running is cluster installed but the RAC feature is disabled.
- Scenario 3: The Oracle home from which the single-instance database is running is not cluster installed.

Use the following procedures to convert your single-instance database on a cluster machine to RAC for all of these scenarios.

Single Instance on a Cluster Running from a Cluster Enabled Oracle Home

Perform the following procedures to convert a single-instance database on a cluster running from a cluster installed (Oracle Database 10g with RAC) Oracle home.

1. Use the DBCA to create a preconfigured image of your single-instance database as described under the heading "[Back up the Original Single-Instance Database](#)" on page D-3. To perform the conversion manually, shut down the single-instance database.
2. To add nodes to your cluster, add and connect these nodes to the cluster as described under the heading "[Perform the Pre-Installation Steps](#)" on page D-3. Ensure that all of these nodes can access the shared storage. Also extend the CRS home to the new nodes using the procedures for "Adding Nodes at the Clusterware Layer on UNIX" as described in the **Oracle Real Application Clusters Administrator's Guide**.
3. From the existing Oracle home, extend this home to the new nodes using the procedure "Adding Nodes at the Oracle RAC Database Layer" as described in the **Oracle Real Application Clusters Administrator's Guide**.
4. From one of the newly added nodes, configure the listeners on the additional nodes using the NetCA. Choose the same port number and protocol that you used on the existing node. If the NetCA displays the existing node in the node list page, then do not select this node because the listener is already configured on it.
5. Convert the database using one of the following procedures:

- [Automated Conversion Procedure](#)
- [Manual Conversion Procedure](#)

Automated Conversion Procedure

1. If you created the preconfigured image of the single instance database as described under the heading "[Back up the Original Single-Instance Database](#)" on page D-3, then use the DBCA to complete the conversion to a RAC database.
2. Start the DBCA from the initial node. Select the names of the nodes that you want to include as part of your cluster database. On the Template Selection page, select the preconfigured template that you created in Step 1 on page D-4. Enter the database name and respond to the remaining DBCA prompts.
3. To use raw devices for the cluster database files, on the Initialization Parameters page enter the raw device name for the SPFILE on the File Locations tab. On the Storage page, replace the default database file names with the raw devices for the control files, redo logs, and datafiles to create the cluster database. Click Finish and create the database.

After creating the RAC database, the DBCA displays the Password Management page on which you must change the passwords for database privileged users who have SYSDBA and SYSOPER roles. Once the DBCA exits, the conversion process is complete.

Manual Conversion Procedure

Because you did not use the DBCA to create a preconfigured image of your single-instance database in Step 1 on page D-4, perform the following steps to complete the conversion:

1. Create the OFA directory structure on all of the nodes that you have added.

See Also: "[UNIX Directory Structures for Real Application Clusters](#)" on page E-2 for more information about OFA

2. If you are converting single-instance database files on a file system to raw devices, then copy the database datafiles, control files, redo logs, and server parameter file to their corresponding raw devices using the `dd` command. Otherwise, continue to the next step.
3. Re-create the control files by executing the `CREATE CONTROLFILE` SQL statement with the `REUSE` keyword and specify `MAXINSTANCES` and `MAXLOGFILES`, and so on, as needed for your RAC configuration. The `MAXINSTANCES` recommended default is 32.
4. Shut down the database instance.
5. If your single-instance database was using an SPFILE parameter file, then create a temporary PFILE from the SPFILE using the following SQL statement:

```
CREATE PFILE='pfile_name' from spfile='spfile_name'
```

6. Set the `CLUSTER_DATABASE` parameter to `TRUE`, set the `INSTANCE_NUMBER` parameter to a unique value for each instance, using a `sid.parameter=value` syntax.

If you optimized memory usage on your single-instance database, adjust the size of the SGA to avoid swapping and paging when you convert to RAC. This is because RAC requires about 350 bytes for each buffer to accommodate the Global Cache Service (GCS). For example, if you have 10,000 buffers, RAC requires about

350*10,000 bytes more memory. Therefore, adjust the size of the SGA by changing the `DB_CACHE_SIZE` and `DB_nK_CACHE_SIZE` parameters accordingly.

7. Start up the database instance using the PFILE created in step 5.
8. If your single-instance database was using automatic undo management, then create an undo tablespace for each additional instance using the `CREATE UNDO TABLESPACE` SQL statement. If you are using raw devices, then ensure that the datafile for the undo tablespace is on the raw device.
9. Create redo threads that have at least two redo logs for each additional instance. If you are using raw devices, then ensure that the redo log files are on raw devices. Enable the new redo threads by using an `ALTER DATABASE` SQL statement. Then shutdown the database instance.
10. Copy the Oracle password file from the initial node, or from the node from which you are working, to the corresponding location on the additional nodes on which the cluster database will have an instance. Make sure that you replace the `ORACLE_SID` name in each password file appropriately for each additional instance.
11. Add `REMOTE_LISTENER=LISTENERS_DB_NAME` and `sid.LOCAL_LISTENER=LISTENER_SID` parameters to the PFILE.
12. Configure the net service entries for the database and instances and address entries for the `LOCAL_LISTENER` for each instance and `REMOTE_LISTENER` in the `tnsnames.ora` file and copy it to all nodes.
13. Create the SPFILE from the PFILE using the procedures under the heading "[Procedures for Migrating to the Server Parameter File](#)" on page 11-3. If you are not using a cluster file system, then ensure that the SPFILE is on a raw device.
14. Create the `$ORACLE_HOME/dbs/initSID.ora` file that contains the following entry:


```
spfile='spfile_path name'
```

Where `spfile_path name` is the complete path name of the SPFILE.
15. Add the configuration for the RAC database and its instance-to-node mapping using `SRVCTL`.
16. Start the RAC database using `SRVCTL`.

After starting the database with `SRVCTL`, your conversion process is complete and, for example, you can execute the following SQL statement to see the statuses of all the instances in your RAC database:

```
select * from v$instance
```

Single Instance on a Cluster Running from a RAC-Disabled Oracle Home

This installation is possible if you performed a one-node cluster (with RAC) installation but later disabled the RAC feature by unlinking it from the `oracle` binary before creating the single instance database. (However, you can also select the "local", "non-cluster" selection on the Node Selection Page to create a non-RAC-enabled single-instance home on a cluster.) Perform the following procedures to convert this type of single-instance database to a RAC database:

1. On the cluster node where the single-instance database is running, execute step 1 of "[Single Instance on a Cluster Running from a Cluster Enabled Oracle Home](#)" on page D-4.

2. Change the directory to the `lib` subdirectory in the `rdbms` directory under the Oracle home.
3. Relink the `oracle` binary by executing the following commands:

```
make -f ins_rdbms.mk rac_on
make -f ins_rdbms.mk ioracle.
```
4. Continue with Step 2 on page D-4.

Single Instance on a Cluster Running from non-Cluster Installed Oracle Home

This installation is only possible if you selected the local installation option on the OUI Specify Hardware Cluster Installation page during an Oracle Database 10g installation.

To convert this database to a RAC database, perform the procedures described under the following headings:

1. ["Back up the Original Single-Instance Database"](#) on page D-3.
2. ["Perform the Pre-Installation Steps"](#) on page D-3.
3. ["Set up the Cluster"](#) on page D-3.
4. ["Install Oracle Database 10g Software with Real Application Clusters"](#). In this step, make sure that you select a new Oracle home other than the one from which the single-instance database was running.

Post-Conversion Steps

After completing the conversion, note the following points as described in the RAC documentation:

- Follow the recommendations for using load balancing and TAF as described in the **Oracle Real Application Clusters Administrator's Guide**
- Use locally managed tablespaces instead of dictionary managed tablespaces to reduce contention and manage sequences in RAC as described in the **Oracle Real Application Clusters Deployment and Performance Guide**
- Follow the guidelines for configuring an interconnect, for using automatic segment space management and for using SRVCTL to administer multiple instances as described in the *Oracle Real Application Clusters Administrator's Guide*

The buffer cache and shared pool capacity requirements in RAC are slightly greater than those in single-instance Oracle databases. Therefore, you may wish to increase the size of the buffer cache by about 10% and the size of the shared pool by about 15%.

Directory Structure for Oracle Database 10g Real Application Clusters Environments

This appendix describes the directory structures for Real Application Clusters (RAC) software environments. The topics in this appendix are:

- [Understanding the Real Application Clusters Directory Structure](#)
- [UNIX Directory Structures for Real Application Clusters](#)

Understanding the Real Application Clusters Directory Structure

When you install Oracle Database 10g with RAC, all subdirectories are under a top-level `ORACLE_BASE`. The `ORACLE_HOME` and `admin` directories are also located under `ORACLE_BASE`.

UNIX Directory Structures for Real Application Clusters

Table E-1 shows the hierarchical directory tree of a sample OFA-compliant database for RAC:

Table E-1 Directory Structure for A Sample OFA-Compliant Environment

Root	Second-Level	Third-Level	Fourth-Level	Fifth-Level
<code>\$ORACLE_BASE</code>				<code>/u01/app/oracle</code> The default <code>ORACLE_BASE</code> directory.
	<code>ORACLE_HOME</code>			<code>/product/10.1</code> The default name of the Oracle home.
	<code>/admin</code>			Administrative directory.
		<code>/db_unique_name</code>		Database unique name, the same as <code>dbname</code> when database name is eight or fewer characters in length.
			<code>/bdump</code> <code>/cdump</code> <code>/hdump</code> <code>/pfile</code> <code>/udump</code>	Dump destinations for database server.
	<code>CRS Home</code>			<code>/crs/10.1</code> The default name of the CRS home.
		<code>/bin</code>		Subtree for Oracle binaries.
		<code>/network</code>		Subtree for Oracle Net.

See Also: *Oracle Database 10g Administrator's Reference Release 1 (10.0) for UNIX Systems: AIX-Based Systems, HP-UX, HP Tru64 UNIX, Linux, and Solaris Operating System (SPARC)* for further information about the `$ORACLE_HOME` and `/admin` directories

Numerics

64-bit

checking system architecture, 2-4, 3-4, 5-4, 6-4

A

Adaptive Server

See Sybase Adaptive Server Enterprise

Advanced Security Option

See Oracle Advanced Security

AdvFS

identifying AdvFS file systems on Tru64 UNIX, 4-30

aio_task_max_num subsystem attribute

recommended value on Tru64 UNIX, 4-15

AIX

activating volume groups, 2-37
APAR download location, 2-8
character raw device names, 2-30
checking disk availability for ASM, 2-30
checking disk availability for raw devices, 2-32
checking existence of nobody user, 2-13
checking maintenance level, 2-7
checking patch requirements, 2-8
checking software requirements, 2-7
checking system architecture, 2-4
checking version, 2-7
configuring disks for ASM, 2-30
configuring new disks, 2-30, 2-32
configuring raw devices, 2-31
creating a raw device mapping file, 2-37
creating a volume group, 2-32
creating identical users, 2-14
creating logical volumes, 2-35
creating raw logical volumes, 2-34
creating volume groups, 2-33
identifying disks for ASM, 2-30
identifying disks for LVM, 2-32
identifying LVM disks, 2-30, 2-32
importing raw device disk group, 2-36
initializing disks for LVM, 2-30, 2-33
location of oratab file, 2-19
patch requirements, 2-7
PTF download location, 2-9
raw device sizes, 2-34

software requirements, 2-6

specifying owner and permissions of raw devices, 2-36, 2-37

AIXwindows

Motif library requirement, 2-6

APAR

checking, 2-8

download location, 2-8

requirements on AIX, 2-7

architecture

checking system architecture, 2-4, 3-4, 5-4, 6-4
Optimal Flexible Architecture (OFA), 8-3

archive logs

destinations, conversion to multi-instance, D-2

ASM

benefits of, 1-5

block and character device names

on HP-UX, 3-27

on Solaris, 6-34

on Tru64 UNIX, 4-30

block device names on Linux, 5-33, 5-35

changing owner and permissions of disks

on AIX, 2-30

on HP-UX, 3-27

on Solaris, 6-34

on Tru64 UNIX, 4-30

changing owner and permissions of raw devices

on Linux, 5-35, 5-36, 5-40

character raw device names on AIX, 2-30

characteristics of failure groups, 1-6, 2-28, 3-25, 4-27, 5-29, 6-31

checking disk availability on AIX, 2-30

checking disk availability on HP-UX, 3-27

checking disk availability on Linux, 5-33, 5-35

checking disk availability on Solaris, 6-33

checking disk availability on Tru64 UNIX, 4-29

compared to logical volume managers, 1-5

compared to RAID, 1-5

configuring disks

on AIX, 2-30

on HP-UX, 3-27

on Linux, 5-31

on Solaris, 6-33

on Tru64 UNIX, 4-29

disk groups, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30

displaying attached disks

- on Linux, 5-33, 5-35
- failure groups, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30
 - examples, 2-28, 3-25, 4-28, 5-29, 6-32
 - identifying, 2-28, 3-25, 4-28, 5-29, 6-32
- identifying available disks on AIX, 2-30
- identifying available disks on HP-UX, 3-27
- identifying available disks on Linux, 5-33, 5-35
- identifying available disks on Solaris, 6-33
- identifying available disks on Tru64 UNIX, 4-30
- identifying disks on AIX, 2-30
- identifying disks on HP-UX, 3-27
- identifying disks on Linux, 5-33, 5-35
- identifying disks on Solaris, 6-33
- identifying disks on Tru64 UNIX, 4-30
- overview, 1-5
- recommendations for disk groups, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30
- redundancy levels, 1-6
- restrictions on use for OCR or voting disk, 2-26, 3-23, 4-25, 5-27, 6-29
- space required for preconfigured database, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
- storage option for datafiles, 2-22, 3-22, 4-21, 5-23, 6-24
- storage option for RAC, 2-22, 3-23, 4-22, 5-24, 6-26

authorized problem analysis report
See APAR

automatic undo management, 12-3

Available
service configuration policy, 9-2

B

- backups
 - conversion to multi-instance, D-2
- base directory
See Oracle base directory
- Bash shell
 - default user startup file, 2-41, 3-38, 4-33, 5-44, 6-42
 - setting shell limits on Linux, 5-18
- .bash_profile file, 2-41, 3-38, 4-33, 5-44, 6-42
- Basic
 - TAF failover policy, 9-3
- block device
 - device name on HP-UX, 3-27, 3-30
 - device name on Linux, 5-33, 5-35
 - device name on Solaris, 6-34, 6-35
 - device name on Tru64 UNIX, 4-30
- boot.local file, 5-17
- Bourne shell
 - default user startup file, 2-41, 3-38, 4-33, 5-44, 6-42
 - setting shell limits on Linux, 5-18

C

- C compiler
 - requirement on AIX, 2-7

- requirement on HP-UX, 3-6
- requirement on Linux, 5-7
- requirement on Tru64 UNIX, 4-6
- See also* Pro*C/C++

C shell

- default user startup file, 2-41, 3-38, 4-33, 5-44, 6-42
- setting shell limits on Linux, 5-18

cfgmgr command, 2-30, 2-32

character device

- device name on HP-UX, 3-27, 3-30
- device name on Solaris, 6-34, 6-35
- device name on Tru64 UNIX, 4-30

character raw device

- device name on AIX, 2-30

chdev command, 2-30, 2-33

chmod command, 2-21, 2-25, 2-30, 3-21, 3-27, 4-21, 4-25, 4-30, 5-22, 5-27, 5-35, 5-36, 5-40, 6-24, 6-29, 6-34

chown command, 2-21, 2-25, 2-30, 3-21, 3-27, 4-21, 4-25, 4-30, 5-22, 5-27, 5-35, 5-36, 5-40, 6-24, 6-29, 6-34

clac_stat command, 3-36

clu_get_info command, 4-31

cluster database

- reasons not to convert to, D-2

cluster file system

- choosing for datafile location, 2-22, 4-22, 5-24, 6-26
- choosing for software location, 2-22, 4-22, 5-24, 6-26
- loading OCFS on Linux, 5-42
- OCFS on Linux, 5-7
- requirement for RAC, 2-23, 2-24, 4-22, 4-23, 4-24, 5-24, 5-25, 6-26, 6-27, 6-28
- storage option for datafiles, 2-22, 3-22, 4-21, 5-23, 6-24
- storage option for RAC, 2-22, 3-23, 4-22, 5-24, 6-26

cluster interconnect

- Hyper Messaging protocol on HP-UX, 3-6

cluster nodes

- activating volume groups
 - on AIX, 2-37
- importing raw device disk group
 - on AIX, 2-36
 - on HP-UX, 3-33
 - on Solaris, 6-38
- specifying uids and gids, 2-13, 3-13, 4-12, 5-13, 6-14
- starting on HP-UX, 3-36

cluster software

- checking on HP-UX, 3-36
- checking on Linux, 5-41
- checking on Solaris, 6-40
- starting Serviceguard on HP-UX, 3-36

clusterware

- See* cluster software
- See* cluster software directories

cmruncl command, 3-36

- cmrunnode command, 3-36
- cmviewcl command, 3-36
- COBDIR environment variable
 - description and setting, 2-43, 3-39, 5-46, 6-44
- components
 - created when using the DBCA, 12-2
- configuration types
 - Customized, 8-2
 - Data Warehouse, 8-2
 - General Purpose, 8-2
 - New Database, 8-2
 - Software Only, 8-2
 - Transaction Processing, 8-2
- configuring
 - global database name, 9-4
 - SID Prefix, 8-6
- connection load balancing, 12-4
- control files, 1-8
 - described, 12-3
 - raw devices for
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- converting
 - from single-instance to Real Application Clusters, B-1, D-1
 - to Real Application Clusters from single-instance Oracle databases, B-1, D-1
- corrective service diskette
 - See* CSD
- Create Database
 - option on Creation Options page, 9-8
- creating
 - Real Application Clusters database with the Database Configuration Assistant, 8-3, 9-2, 9-3
- cross-node registration, 12-4
- CRS
 - OCR restriction, 2-26, 3-23, 4-25, 5-27, 6-29
 - raw device for OCR on AIX, 2-35
 - raw device for OCR on HP-UX, 3-32
 - raw device for OCR on Linux, 5-38
 - raw device for OCR on Solaris, 6-37
 - raw device for voting disk on AIX, 2-35
 - raw device for voting disk on HP-UX, 3-32
 - raw device for voting disk on Linux, 5-38
 - voting disk restriction, 2-26, 3-23, 4-25, 5-27, 6-29
- CSD
 - download location for MQSeries or WebSphere MQ
 - on HP-UX, 3-9
 - on Tru64 UNIX, 4-8
 - download location for WebSphere MQ
 - on AIX, 2-9
 - on Linux, 5-8
 - on Solaris, 6-10
 - requirements
 - on AIX, 2-7, 2-8
 - on HP-UX, 3-8

- on Linux, 5-8
 - on Solaris, 6-10
 - on Tru64 UNIX, 4-8
- csh.login file, 5-18
- csh.login.local file, 5-18
- custom database
 - failure groups for ASM, 2-28, 3-25, 4-28, 5-29, 6-32
 - requirements when using ASM, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
- Custom installation type
 - reasons for choosing, 2-9, 3-10, 4-8, 5-9, 6-10
- Customized
 - configuration type, 8-2, 8-3

D

- data loss
 - minimizing with ASM, 1-6, 2-28, 3-25, 4-27, 5-29, 6-31
- Data Warehouse
 - configuration type, 8-2
- database
 - ASM requirements, 1-7
 - components, created when using the DBCA, 12-2
 - configurations, types, 8-2
- Database Configuration Assistant
 - components created by, 12-2
 - control files, 12-3
 - creating Real Application Clusters database
 - after installation, 9-3
 - during installation, 8-3, 9-2
 - Database Identification page, 9-4
 - Database Templates page, 9-3
 - datafiles, 12-2
 - deleting databases with, 9-9
 - deleting Real Application Clusters databases, 9-9
 - initialization parameter files, 12-4
 - List of Cluster Databases page, 9-9
 - Operations page, 9-3, 9-9
 - redo log files, 12-3
 - rollback segments, 12-3
 - Summary dialog, 9-9
 - tablespaces, 12-2
 - using, 9-2
 - Welcome page, 9-3
- database configuration type
 - selecting, 8-2
- database configuration types, 8-2
- Database Identification page, 9-4
- Database Templates page, 9-3
- databases
 - ASM requirements, 2-27, 3-24, 4-27, 5-29, 6-31
- datafiles, 1-8
 - and the DBCA, 12-3
 - creating separate directories for, 2-25, 4-24, 5-26, 6-28
 - described, 12-2
 - managing with ASM, 1-5
 - minimum disk space for, 2-24, 4-23, 5-25, 6-27

- options for placing on file system, 2-23, 4-22, 5-24, 6-26
- recommendations for file system, 2-23, 4-23, 5-25, 6-27
- setting permissions on datafile directories, 2-25, 4-25, 5-27, 6-29
- storage options, 2-22, 3-22, 4-21, 5-23, 6-24

dba group

- and ASM disks on AIX, 2-30
- and ASM disks on HP-UX, 3-27
- and ASM disks on Linux, 5-35, 5-36, 5-40
- and ASM disks on Solaris, 6-34
- and ASM disks on Tru64 UNIX, 4-30
- and SYSDBA privilege, 2-9, 3-10, 4-8, 5-9, 6-10
- creating, 2-11, 3-11, 4-10, 5-11, 6-12
- creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
- description, 2-9, 3-10, 4-8, 5-9, 6-10
- granting privileges on HP-UX, 3-15
- raw device group
 - on AIX, 2-36, 2-37
 - on HP-UX, 3-33
 - on Linux, 5-39
 - on Solaris, 6-38

DBCA_RAW_CONFIG environment variable

- description and setting, 2-43, 3-39, 5-46, 6-44
- on AIX, 2-39
- on HP-UX, 3-35
- on Linux, 5-41
- on Solaris, 6-40

DCE

- patches required on HP-UX, 3-8
- requirements for ASO on HP-UX, 3-7
- requirements for ASO on Tru64 UNIX, 4-6
- requirements on HP-UX, 3-7
- requirements on Tru64 UNIX, 4-6

dedicated servers, 12-4

default file mode creation mask

- setting, 2-41, 3-37, 4-32, 5-44, 6-42

device names

- for raw devices on Linux, 5-39
- IDE disks on Linux, 5-33, 5-35, 5-37
- on AIX, 2-30
- on HP-UX, 3-27, 3-30
- on Linux, 5-33, 5-35
- on Solaris, 6-34, 6-35
- on Tru64 UNIX, 4-30
- RAID on Linux, 5-33, 5-35, 5-37
- raw devices on Linux, 5-35, 5-36
- SCSI disks on Linux, 5-33, 5-35, 5-37

device numbers

- identifying major numbers on AIX, 2-33
- identifying minor numbers on HP-UX, 3-31

directory

- creating separate datafile directories, 2-25, 4-24, 5-26, 6-28
- database file directory, 2-23, 4-23, 5-25, 6-27
- Oracle base directory, 2-17, 3-18, 4-17, 5-19, 6-19
- Oracle home directory, 2-18, 3-19, 4-18, 5-20, 6-21
- Oracle Inventory directory, 2-18, 3-19, 4-17, 5-19, 6-20
- oraInventory, 2-18, 3-19, 4-17, 5-19, 6-20
- permission for datafile directories, 2-25, 4-25, 5-27, 6-29

directory structure, E-1

- UNIX, E-2

disk

- changing permissions and owner for ASM
 - on HP-UX, 3-27

disk devices

- managing with ASM, 1-5

disk group

- ASM, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30
- recommendations for ASM disk groups, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30

disk space

- checking, 2-4, 3-4, 4-3, 5-3, 6-3
- requirement for Oracle base directory, 2-20, 2-21, 4-20, 5-21, 5-22, 6-23
- requirements for preconfigured database in ASM, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31

disklabel command, 4-30

disks

- changing permissions and owner for ASM
 - on AIX, 2-30
 - on HP-UX, 3-27
 - on Solaris, 6-34
 - on Tru64 UNIX, 4-30
- checking availability for ASM on AIX, 2-30
- checking availability for ASM on HP-UX, 3-27
- checking availability for ASM on Linux, 5-33, 5-35
- checking availability for ASM on Solaris, 6-33
- checking availability for ASM on Tru64 UNIX, 4-29
- checking availability for raw devices on AIX, 2-32
- checking availability for raw devices on HP-UX, 3-30
- checking availability for raw devices on Solaris, 6-35
- configuring for ASM on AIX, 2-30
- configuring for ASM on HP-UX, 3-27
- configuring for ASM on Linux, 5-31
- configuring for ASM on Solaris, 6-33
- configuring for ASM on Tru64 UNIX, 4-29
- configuring new disks, 2-30, 2-32
- configuring raw devices on AIX, 2-31
- configuring raw devices on HP-UX, 3-28
- configuring raw devices on Solaris, 6-34
- device names on Solaris, 6-35
- displaying attached disks
 - on Linux, 5-33, 5-35, 5-37
- identifying AdvFS file systems on Tru64 UNIX, 4-30
- identifying file systems on Solaris, 6-34, 6-35
- identifying LSM disks on Tru64 UNIX, 4-30
- identifying LVM disks on AIX, 2-30, 2-32
- identifying LVM disks on HP-UX, 3-27, 3-30
- identifying LVM disks on Solaris, 6-34, 6-35
- identifying UFS file systems on Tru64 UNIX, 4-30

- identifying unused disks on Tru64 UNIX, 4-30
- raw voting disk on AIX, 2-35
- raw voting disk on HP-UX, 3-32
- raw voting disk on Linux, 5-38
- raw voting disk on Solaris, 6-37
- DISPLAY environment variable, 3-17
 - setting, 2-41, 2-42, 3-37, 3-38, 4-32, 4-33, 5-44, 5-45, 6-42
- Distributed Computing Environment
 - See DCE
- dupatch command, 4-8
- Dynamic Server
 - See Informix Dynamic Server

E

- emulator
 - installing from X emulator, 2-2, 3-2, 4-2, 5-2, 6-2
- env command, 2-43, 3-40, 4-34, 5-46, 6-44
- environment
 - checking settings, 2-43, 3-40, 4-34, 5-46, 6-44
 - configuring for oracle user, 2-41, 3-37, 4-32, 5-44, 6-42
- environment variables
 - COBDIR, 2-43, 3-39, 5-46, 6-44
 - DBCA_RAW_CONFIG, 2-39, 2-43, 3-35, 3-39, 5-41, 5-46, 6-40, 6-44
 - DISPLAY, 2-41, 2-42, 3-17, 3-37, 3-38, 4-32, 4-33, 5-44, 5-45, 6-42
 - ORACLE_BASE, 2-18, 2-21, 2-41, 3-18, 3-22, 3-37, 4-17, 4-21, 4-32, 5-19, 5-23, 5-44, 6-20, 6-24, 6-42
 - ORACLE_HOME, 2-41, 2-43, 3-37, 3-39, 4-32, 4-34, 5-43, 5-44, 5-46, 6-41, 6-42, 6-44
 - ORACLE_SID, 2-41, 3-37, 4-32, 5-44, 6-42
 - PATH, 2-41, 2-43, 3-37, 3-39, 4-32, 5-44, 5-46, 6-42, 6-44
 - removing from shell startup file, 2-42, 3-38, 4-33, 5-44, 6-42
 - required for Pro*COBOL, 2-43, 3-39, 5-46, 6-44
 - SHELL, 2-41, 3-38, 4-33, 5-44, 6-42
 - TEMP and TMPDIR, 2-3, 2-42, 3-3, 3-38, 4-3, 4-33, 5-3, 5-45, 6-3, 6-43
 - TNS_ADMIN, 2-43, 3-39, 4-34, 5-46, 6-44
- equivalence
 - setting up UNIX user equivalence, 2-15, 3-15, 4-14, 6-17
- errata
 - checking glibc errata on Red Hat, 5-8
 - Red Hat
 - glibc errata, 5-8
 - Red Hat kernel errata, 5-8
 - download location, 5-8
- ESQL/C
 - See Informix ESQL/C
 - /etc/csh.login file, 5-18
 - /etc/csh.login.local file, 5-18
 - /etc/hosts.equiv file, 2-15, 3-15, 4-14, 6-17
 - /etc/init.d/boot.local file, 5-17
 - /etc/lvmrc file, 3-29
 - /etc/pam.d/login file, 5-18

- /etc/privgroup file, 3-15
- /etc/profile file, 5-18
- /etc/profile.local file, 5-18
- /etc/raw file, 5-39
- /etc/security/limits.so file, 5-18
- /etc/sysconfig/rawdevices file, 5-35, 5-36, 5-39
- /etc/sysconfigtab file, 4-16
- /etc/sysctl.conf file, 5-17
- /etc/system file, 6-19
- EXAMPLE tablespace
 - raw device for
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36
- examples
 - ASM failure groups, 2-28, 3-25, 4-28, 5-29, 6-32
 - Oracle base directories, 2-17, 3-18, 4-17, 5-19, 6-20
- external jobs
 - UNIX user required for, 2-9, 3-10, 4-9, 5-9, 6-10
- external redundancy
 - ASM redundancy level, 1-6
- extjob executable
 - UNIX user required for, 2-9, 3-10, 4-9, 5-9, 6-10
- extjob user
 - creating on HP-UX, 3-13

F

- failover
 - and service registration, 12-4
- failure group
 - ASM, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30
 - characteristics of ASM failure group, 1-6, 2-28, 3-25, 4-27, 5-29, 6-31
 - examples of ASM failure groups, 2-28, 3-25, 4-28, 5-29, 6-32
- fdisk command, 5-33, 5-35, 5-37, 5-38
- features, new, 4-xxiii
- fifo_do_adaptive subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- file command, 4-29
- file mode creation mask
 - setting, 2-41, 3-37, 4-32, 5-44, 6-42
- file system
 - appropriate for Oracle base directory, 2-21, 3-21, 4-20, 5-22, 6-23
 - datafile and recovery file placement
 - options, 2-23, 4-22, 5-24, 6-26
 - identifying AdvFS file systems on Tru64 UNIX, 4-30
 - identifying formatted devices on Solaris, 6-34, 6-35
 - identifying UFS file systems on Tru64 UNIX, 4-30
 - requirements for RAC, 2-23, 2-24, 4-22, 4-23, 4-24, 5-24, 5-25, 6-26, 6-27, 6-28
 - storage option for datafiles, 2-22, 3-22, 4-21, 5-23, 6-24
 - using for datafiles, 2-23, 4-23, 5-25, 6-27
- file-max file, 5-16

- file-max parameter
 - recommended value on Linux, 5-16
- files
 - .bash_profile, 2-41, 3-38, 4-33, 5-44, 6-42
 - boot.local, 5-17
 - control files
 - raw devices for on AIX, 2-35
 - raw devices for on HP-UX, 3-32
 - raw devices for on Linux, 5-38
 - raw devices for on Solaris, 6-37
 - default shell startup file, 5-18
 - editing shell startup file, 2-41, 3-38, 4-33, 5-44, 6-42
 - /etc/csh.login, 5-18
 - /etc/csh.login.local, 5-18
 - /etc/hosts.equiv, 2-15, 3-15, 4-14, 6-17
 - /etc/init.d/boot.local, 5-17
 - /etc/lvmrc, 3-29
 - /etc/pam.d/login, 5-18
 - /etc/privgroup, 3-15
 - /etc/profile, 5-18
 - /etc/profile.local, 5-18
 - /etc/raw, 5-39
 - /etc/security/limits.so, 5-18
 - /etc/sysconfig/rawdevices, 5-35, 5-36, 5-39
 - /etc/sysconfigtab, 4-16
 - /etc/sysctl.conf, 5-17
 - /etc/system, 6-19
 - .login, 2-41, 3-38, 4-33, 5-44, 6-42
 - oraInst.loc, 2-19, 3-20, 4-19, 5-10, 5-21, 6-22
 - oratab, 2-19, 3-20, 4-19, 5-21, 6-22
 - password file
 - raw device for on AIX, 2-35
 - raw device for on HP-UX, 3-32
 - raw device for on Linux, 5-38
 - raw device for on Solaris, 6-37
 - /proc/sys/fs/file-max, 5-16
 - /proc/sys/kernel/sem, 5-16
 - /proc/sys/kernel/shmall, 5-16
 - /proc/sys/kernel/shmmax, 5-16
 - /proc/sys/kernel/shmmni, 5-16
 - /proc/sys/net/ipv4/ip_local_port_range, 5-16
 - .profile, 2-41, 3-38, 4-33, 5-44, 6-42
 - profile.local, 5-18
 - raw device mapping file, 2-43, 3-39, 5-46, 6-44
 - on AIX, 2-37
 - on HP-UX, 3-34
 - on Linux, 5-40
 - on Solaris, 6-38
 - redo log files
 - raw devices for on AIX, 2-35
 - raw devices for on HP-UX, 3-32
 - raw devices for on Linux, 5-38
 - raw devices for on Solaris, 6-37
 - .rhosts, 2-15, 3-15, 4-14, 6-17
 - server parameter file
 - raw devices for on AIX, 2-35
 - raw devices for on HP-UX, 3-32
 - raw devices for on Linux, 5-38
 - raw devices for on Solaris, 6-37

- SPFILE
 - raw device for on AIX, 2-35
 - raw devices for on HP-UX, 3-32
 - raw devices for on Linux, 5-38
 - raw devices for on Solaris, 6-37

- filesets
 - AIX requirements, 2-6
 - checking on AIX, 2-7
- for raw devices on Linux, 5-39
- format command, 6-33, 6-35

G

- Gateway
 - See* Oracle Messaging Gateway
- General Purpose
 - configuration type, 8-2
 - description of configuration type, 8-2
- Generate Database Creation Scripts
 - on Creation Options page, 9-9
- getconf command, 2-4, 3-4, 5-4, 6-4
- getprivgrp command, 3-16
- gh_chunks subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- gid
 - identifying existing, 2-14, 3-13, 4-12, 5-13, 6-14
 - specifying, 2-14, 3-14, 4-13, 5-13, 6-14
 - specifying on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
- glibc
 - Red Hat errata, 5-8
- global database name, 9-4
- group, 2-14
- group IDs
 - identifying existing, 2-14, 3-13, 4-12, 5-13, 6-14
 - specifying, 2-14, 3-14, 4-13, 5-13, 6-14
 - specifying on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
- groupadd command, 4-13, 5-11
- groups
 - checking for existing oinstall group, 2-10, 2-14, 3-11, 4-10, 5-10, 6-12
 - creating identical groups on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - creating the dba group, 2-11, 3-11, 4-10, 5-11, 6-12
 - creating the oinstall group, 2-10, 3-11, 4-10, 5-10, 6-11
 - creating the oper group, 2-11, 3-12, 4-10, 5-11, 6-12
 - specifying when creating UNIX users, 2-14, 3-14, 4-13, 5-13, 6-14
 - UNIX OSDBA group (dba), 2-9, 3-10, 4-8, 5-9, 6-10
 - UNIX OSOPER group (oper), 2-9, 3-10, 4-8, 5-9, 6-10

H

- hardware requirements, 2-3, 3-3, 4-3, 5-3, 6-3
- high redundancy

- ASM redundancy level, 1-6
- home directory
 - See Oracle home directory
- hosts.equiv file, 2-15, 3-15, 4-14, 6-17
- HP 9000 system
 - HP-UX operating system requirement, 3-6
- HP-UX
 - block and character device names, 3-27, 3-30
 - checking disk availability for ASM, 3-27
 - checking disk availability for raw devices, 3-30
 - checking HyperFabric software, 3-36
 - checking patch requirements, 3-8
 - checking quality pack requirement, 3-8
 - checking Serviceguard, 3-36
 - checking software requirements, 3-7
 - checking system architecture, 3-4
 - checking version, 3-7
 - configuring disks for ASM, 3-27
 - configuring kernel parameters, 3-16
 - configuring raw devices, 3-28
 - creating a raw device mapping file, 3-34
 - creating a volume group, 3-30
 - creating a volume group directory on HP-UX
 - volume group
 - creating a directory for on HP-UX, 3-31
 - creating an unprivileged user, 3-13
 - creating logical volumes, 3-32
 - creating raw logical volumes, 3-31
 - creating required X library symbolic links, 3-9
 - creating volume groups, 3-31
 - disabling volume group activation, 3-29
 - granting privileges to OSDBA group, 3-15
 - identifying disks for ASM, 3-27
 - identifying disks for LVM, 3-30
 - identifying LVM disks, 3-27, 3-30
 - importing raw device disk group, 3-33
 - initializing disks for LVM, 3-31
 - location of oratab file, 3-20
 - maxvgs parameter, 3-31
 - patch download location, 3-9
 - patch requirements, 3-7
 - quality pack download location, 3-8
 - quality pack requirements, 3-8
 - raw device sizes, 3-31
 - requirements on HP 9000 systems, 3-6
 - setting kernel parameters, 3-17
 - software requirements, 3-6
 - specifying owner and permissions of raw devices, 3-33
 - starting a cluster node, 3-36
 - System Administration Manager (SAM), 3-17
 - verifying that cluster software is running, 3-36
- HP-UX Developer's Toolkit
 - Motif requirement, 3-6
- Hyper Messaging Protocol
 - using as a cluster interconnect, 3-6
- HyperFabric software
 - checking on HP-UX, 3-36
 - HP-UX requirement, 3-6
- IBM MQSeries
 - requirement on Tru64 UNIX, 4-6
- IBM WebSphere MQ
 - requirement on AIX, 2-6
 - requirement on HP-UX, 3-6
 - requirement on Linux, 5-6
 - requirement on Solaris, 6-7
 - requirement on Tru64 UNIX, 4-6
- id command, 2-13, 2-14, 3-13, 4-12, 5-12, 5-13, 6-14
- IDE disks
 - device names on Linux, 5-33, 5-35, 5-37
- Imake
 - requirement on HP-UX, 3-6
- importvg command, 2-36
- Informix Dynamic Server
 - requirements on HP-UX, 3-7
 - requirements on Solaris, 6-8
- Informix ESQL/C
 - requirements on HP-UX, 3-7
 - requirements on Solaris, 6-8
- Ingres II
 - requirements on HP-UX, 3-7
 - requirements on Solaris, 6-8
- initialization parameter files, 12-4
 - for instances, 11-2
 - listener parameters, 12-5
- initialization parameters
 - DISPATCHERS, 12-4
 - LOCAL_LISTENER, 12-4
 - REMOTE_LISTENER, 12-4, 12-8
- insmod command, 5-42
- installation
 - directory structure, E-1
 - ldap.ora file, 12-6
 - listener.ora file, 12-5
 - overview, 1-4
 - requirements, hardware, 1-2
 - requirements, software, 1-2
 - tnsnames.ora file, 12-7
 - verifying raw devices, 9-3
- installation types
 - and ASM requirements, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
- instance
 - instance identifier (SID), 2-41, 3-37, 4-32, 5-44, 6-42
 - preferences, 9-2
- instances
 - initialization parameter files, 11-2
 - SID Prefix, 8-6
- instfix command, 2-8
- ioscan command, 3-27, 3-30
- ip_local_port_range file, 5-16
- ip_local_port_range parameter
 - recommended value on Linux, 5-16
- ipc
 - kernel subsystem on Tru64 UNIX, 4-15
- isainfo command, 2-4, 3-4, 5-4, 6-4

J

Java

font package requirements for Solaris, 6-7

JDK

AIX requirements, 2-6

font packages required on Solaris, 6-7

HP-UX requirements, 3-6

Tru64 UNIX requirements, 4-5

K

kernel

Red Hat errata, 5-8

download location, 5-8

kernel parameters

checking on Linux, 5-17

checking on Solaris, 6-19

checking on Tru64 UNIX, 4-16

configuring on HP-UX, 3-16

configuring on Linux, 5-16

configuring on Solaris, 6-18

making changes persist on Linux, 5-17

making changes persist on Solaris, 6-19

making changes persist on Tru64 UNIX, 4-16

maxvgs on HP-UX, 3-31

setting on HP-UX, 3-17

setting on Linux, 5-17

Korn shell

default user startup file, 2-41, 3-38, 4-33, 5-44, 6-42

setting shell limits on Linux, 5-18

ksh

See Korn shell

ksi_alloc_max parameter

recommended value on HP-UX, 3-16

L

ldap.ora file, 12-6

creating, 12-6

default configuration, 12-6

library

Motif library, 2-6

Lightweight Directory Access Protocol (LDAP), 12-6

limit command, 5-18

limits.so file, 5-18

Linux

binding raw devices to partitions, 5-39

block device names, 5-33, 5-35

checking disk availability for ASM, 5-33, 5-35

checking distribution, 5-7

checking existence of nobody user, 5-12

checking kernel parameters, 5-17

checking patch requirements, 5-8

checking raw device bindings, 5-35, 5-36

checking software requirements, 5-7

checking system architecture, 5-4

checking version, 5-7

cluster file system, 5-7

configuring disks for ASM, 5-31

configuring kernel parameters, 5-16

creating a raw device mapping file, 5-40

creating partitions, 5-38

creating raw partitions, 5-37

displaying attached disks, 5-33, 5-35, 5-37

errata required for Red Hat, 5-8

IDE disk device names, 5-33, 5-35, 5-37

identifying disks for ASM, 5-33, 5-35

identifying disks for raw devices, 5-37

identifying partitions, 5-37

loading hangcheck-timer module, 5-42

loading OCFS, 5-42

location of oraInst.loc file, 5-10

location of oratab file, 5-21

making kernel parameter changes persist, 5-17

patch requirements, 5-8

RAID device names, 5-33, 5-35, 5-37

raw device sizes, 5-38

raw devices, 5-35, 5-36

SCSI disk device names, 5-33, 5-35, 5-37

setting kernel parameters, 5-17

setting shell limits, 5-18

software requirements, 5-6

specifying owner and permissions on raw devices, 5-39

verifying that cluster software is running, 5-41

List of Cluster Databases page, 9-9

listener

identifying Oracle home for, 2-40, 3-37, 4-31, 5-43, 6-41

local, 12-4

lsnrctl command, 2-41, 3-37, 4-32, 5-43, 6-41

registration, 12-6

stopping, 2-40, 2-41, 3-36, 3-37, 4-31, 4-32, 5-43, 6-41

stopping existing listener process, 2-40, 3-36, 4-31, 5-43, 6-41

listener.ora file, 12-5

configuring, 12-5

default configuration, 12-5

listeners

listener.ora file, 12-5

load balancing

and service registration, 12-4

load_ocfs command, 5-42

local device

using for datafiles, 2-24, 4-23, 5-25, 6-27

local listener, 12-4

logical volume manager

See LVM

logical volumes

creating on AIX, 2-35

creating on HP-UX, 3-32

creating on Solaris, 6-37

.login file, 2-41, 3-38, 4-33, 5-44, 6-42

login file, 5-18

lsdev command, 2-30, 2-32, 5-33, 5-35

lspp command, 2-7, 2-9

LSM

identifying volume group devices on Tru64

- UNIX, 4-30
- lsmod command, 5-42
- lsnrctl command, 2-41, 3-37, 4-32, 5-43, 6-41
- lspv command, 2-30, 2-32, 2-36
- lvcreate command, 3-32
- LVM
 - checking disk availability on HP-UX, 3-30
 - compared to ASM, 1-5
 - configuring raw devices on AIX, 2-31
 - configuring raw devices on HP-UX, 3-28
 - configuring raw devices on Solaris, 6-34
 - creating a volume group directory on HP-UX, 3-31
 - creating a volume group on AIX, 2-32
 - creating a volume group on HP-UX, 3-30
 - creating a volume group on Solaris, 6-35
 - creating raw logical volumes on AIX, 2-34
 - creating raw logical volumes on HP-UX, 3-31
 - creating raw logical volumes on Solaris, 6-36
 - creating volume groups on AIX, 2-33
 - creating volume groups on HP-UX, 3-31
 - creating volume groups on Solaris, 6-36
 - disabling volume group activation on HP-UX, 3-29
 - identifying available disks on AIX, 2-32
 - identifying available disks on HP-UX, 3-30
 - identifying available disks on Solaris, 6-35
 - identifying LVM disks on HP-UX, 3-30
 - identifying major device numbers on AIX, 2-33
 - identifying minor device numbers on HP-UX, 3-31
 - identifying volume group devices on AIX, 2-30, 2-32
 - identifying volume group devices on HP-UX, 3-27
 - identifying volume group devices on Solaris, 6-34, 6-35
 - initializing disks on AIX, 2-30, 2-33
 - initializing disks on HP-UX, 3-31
 - initializing disks on Solaris, 6-36
 - maximum number of volume groups on HP-UX, 3-31
 - recommendations for ASM, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30
 - Veritas Volume Manager on Solaris, 6-34, 6-35
- lvmmc file, 3-29

M

- maintenance level
 - checking on AIX, 2-7
- major device numbers
 - identifying on AIX, 2-33
- mapping file
 - for raw devices, 2-43, 3-39, 5-46, 6-44
 - on AIX, 2-37
 - on HP-UX, 3-34
 - on Linux, 5-40
 - on Solaris, 6-38
- mask

- setting default file mode creation mask, 2-41, 3-37, 4-32, 5-44, 6-42
- max_async_req subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- max_objs subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- max_per_proc_address_space subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- max_per_proc_data_size subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- max_per_proc_stack_size subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- max_sessions subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- max_thread_proc parameter
 - recommended value on HP-UX, 3-16
- maxdsiz parameter
 - recommended value on HP-UX, 3-16
- maxdsiz_64bit parameter
 - recommended value on HP-UX, 3-16
- maxssiz parameter
 - recommended value on HP-UX, 3-16
- maxssiz_64bit parameter
 - recommended value on HP-UX, 3-16
- maxswapchunks parameter
 - recommended value on HP-UX, 3-16
- maxuprc parameter
 - recommended value on HP-UX, 3-16
- maxvgs parameter, 3-31
- MC/ServiceGuard
 - See* Serviceguard
- memory requirements, 2-3, 3-3, 4-3, 5-3, 6-3
- Messaging Gateway
 - See* Oracle Messaging Gateway
- migration
 - from single-instance, *see* conversion, D-2
- minor device numbers
 - identifying on HP-UX, 3-31
- mkdir command, 2-21, 2-25, 3-21, 4-21, 4-25, 5-22, 5-27, 6-24, 6-29
- mklv command, 2-35
- mknod command, 3-31
- mkvg command, 2-33
- mode
 - setting default file mode creation mask, 2-41, 3-37, 4-32, 5-44, 6-42
- Motif
 - requirement on Tru64 UNIX, 4-6
- mount command, 4-30
- mount point
 - for Oracle base directory, 2-17, 3-18, 4-17, 5-19
- MQSeries
 - checking requirements on Tru64 UNIX, 4-7
 - CSD download location
 - for HP-UX, 3-9
 - for Tru64 UNIX, 4-8
 - CSDs required
 - on AIX, 2-8
 - on HP-UX, 3-8
 - on Linux, 5-8

- on Tru64 UNIX, 4-8
- requirement on Tru64 UNIX, 4-6
- msg_size subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- msgmap parameter
 - recommended value on HP-UX, 3-16
- msgmni parameter
 - recommended value on HP-UX, 3-16
- msgseg parameter
 - recommended value on HP-UX, 3-16
- msgtql parameter
 - recommended value on HP-UX, 3-16

N

- NAS
 - requirement for RAC, 2-23, 2-24, 4-22, 4-23, 4-24, 5-24, 5-25, 6-26, 6-27, 6-28
- NCR Teradata ODBC Driver
 - requirements on HP-UX, 3-7
 - requirements on Solaris, 6-8
- ncsize parameter
 - recommended value on HP-UX, 3-16
- network attached storage
 - See* NAS
- network configuration files
 - ldap.ora.ora, 12-6
 - listener.ora, 12-5
 - sqlnet.ora, 12-11
 - tnsnames.ora, 12-7
- network directory, E-2
- Network Information Services
 - See* NIS
- New Database
 - configuration type, 8-2
- new features, 4-xxiii
- new_wire_method subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- nfile parameter
 - recommended value on HP-UX, 3-16
- nflocks parameter
 - recommended value on HP-UX, 3-16
- ninode parameter
 - recommended value on HP-UX, 3-17
- NIS
 - alternative to local users and groups, 2-10, 3-11, 4-9, 5-10, 6-11
- nkthread parameter
 - recommended value on HP-UX, 3-17
- nobody user
 - checking existence of, 2-13, 4-12, 5-12, 6-14
 - description, 2-9, 4-9, 5-9, 6-10
- node
 - starting a cluster node on HP-UX, 3-36
- nofile
 - shell limit on Linux, 5-18
- None
 - TAF failover policy, 9-3
- normal redundancy
 - ASM redundancy level, 1-6

- Not used
 - service configuration policy, 9-2
- nproc
 - shell limit on Linux, 5-18
- nproc parameter
 - recommended value on HP-UX, 3-17

O

- OCFS
 - checking on Linux, 5-7
 - download location for Linux, 5-7
 - Linux cluster file system, 5-7
 - loading on Linux, 5-42
 - restriction for Oracle base directory on Linux, 5-22
- OCR
 - raw device for
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
 - restriction on using ASM for, 2-26, 3-23, 4-25, 5-27, 6-29
- ODBC driver
 - Teradata requirement on HP-UX, 3-7
 - Teradata requirement on Solaris, 6-8
- OFA
 - recommendations for Oracle base directory, 2-17, 3-18, 4-17, 5-19, 6-19
 - recommended path for Oracle base directory, 2-17, 3-18, 4-17, 5-19, 6-19
 - recommended path for Oracle home directory, 2-18, 3-19, 4-18, 5-20, 6-21
 - recommended path for Oracle Inventory directory, 2-18, 3-19, 4-17, 5-19, 6-20
- oinstall group
 - checking for existing, 2-10, 3-11, 4-10, 5-10, 6-12
 - creating, 2-10, 3-11, 4-10, 5-10, 6-11
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-10, 3-10, 4-9, 5-9, 6-11
- Open Client
 - See* Sybase Open Client
- Open Database Connectivity
 - See* ODBC
- oper group
 - and SYSOPER privilege, 2-9, 3-10, 4-8, 5-9, 6-10
 - creating, 2-11, 3-12, 4-10, 5-11, 6-12
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-9, 3-10, 4-8, 5-9, 6-10
- operating system
 - checking distribution and version of Linux, 5-7
 - checking version of AIX, 2-7
 - checking version of HP-UX, 3-7
 - checking version of Solaris, 6-8
 - checking version of Tru64 UNIX, 4-7
- operating system requirements
 - AIX, 2-6

- HP-UX, 3-6
- Linux, 5-6
- Solaris, 6-6, 6-7
- Tru64 UNIX, 4-5
- Operations page, 9-3, 9-9
- Optimal Flexible Architecture
 - See OFA
- Optimal Flexible Architecture (OFA), 8-3
- Oracle Advanced Security
 - patches required on HP-UX, 3-8
 - requirements on HP-UX, 3-7
 - requirements on Tru64 UNIX, 4-6
- Oracle base directory
 - and ORACLE_BASE environment variable, 2-18, 3-18, 4-17, 5-19, 6-20
 - creating, 2-21, 3-21, 4-21, 5-22, 6-24
 - creating new, 2-21, 3-21, 4-20, 5-22, 6-23
 - description, 2-17, 3-18, 4-17, 5-19, 6-19
 - determining disk space on, 2-20, 4-20, 5-22, 6-23
 - disk space requirements, 2-20, 2-21, 4-20, 5-21, 5-22, 6-23
 - equivalent directory on Windows, 2-17, 3-18, 4-17, 5-19, 6-19
 - examples, 2-17, 3-18, 4-17, 5-19, 6-20
 - identifying appropriate file system, 2-21, 3-21, 4-20, 5-22, 6-23
 - identifying existing, 2-19, 3-20, 5-20, 6-21
 - mount point for, 2-17, 3-18, 4-17, 5-19
 - OCFS restriction on Linux, 5-22
 - recommended path, 2-17, 3-18, 4-17, 5-19, 6-19
 - relationship with Oracle Software Owner user, 2-17, 3-18, 4-17, 5-19, 6-20
 - requirement for, 2-17, 3-18, 4-17, 5-19, 6-19
 - requirement for RAC installations, 2-21, 3-21, 4-20, 5-22, 6-23
 - requirements for existing directory, 2-20, 3-21, 4-19, 5-21, 6-22
- Oracle Client
 - hardware requirements, 2-3, 3-3, 4-3, 5-3, 6-3
- Oracle Cluster File System
 - See OCFS
- Oracle Cluster Ready Services
 - See CRS
- Oracle Cluster Register
 - See OCR
- Oracle Cluster Registry
 - See OCR
- Oracle Database
 - creating datafile directories, 2-25, 4-24, 5-26, 6-28
 - datafile storage options, 2-22, 3-22, 4-21, 5-23, 6-24
 - hardware requirements, 2-3, 3-3, 4-3, 5-3, 6-3
 - minimum disk space requirements, 2-24, 4-23, 5-25, 6-27
 - privileged groups, 2-9, 3-10, 4-8, 5-9, 6-10
 - requirements with ASM, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
 - setting ORACLE_SID environment variable, 2-41, 3-37, 4-32, 5-44, 6-42
- Oracle DCE Integration
 - requirements on HP-UX, 3-7
 - requirements on Tru64 UNIX, 4-6
- Oracle Hangcheck Timer
 - loading module on Linux, 5-42
- Oracle home directory
 - description, 2-18, 3-19, 4-18, 5-20, 6-21
 - identifying for listener, 2-40, 3-37, 4-31, 5-43, 6-41
 - recommended path, 2-18, 3-19, 4-18, 5-20, 6-21
 - requirement for, 2-18, 3-19, 4-18, 5-20, 6-21
 - requirements, 2-18, 3-19, 4-18, 5-20, 6-21
 - using to identify Oracle base directory, 2-19, 3-20, 4-19, 5-21, 6-22
- Oracle home name, 2-18, 3-19, 4-18, 5-20, 6-21
- Oracle Inventory
 - description, 2-18, 3-19, 4-17, 5-19, 6-20
 - pointer file, 5-10
- Oracle Inventory directory
 - description, 2-18, 3-19, 4-17, 5-19, 6-20
 - recommended path, 2-18, 3-19, 4-17, 5-19, 6-20
- Oracle Inventory group
 - checking for existing, 2-10, 3-11, 4-10, 5-10, 6-12
 - creating, 2-10, 3-11, 4-10, 5-10, 6-11
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-10, 3-10, 4-9, 5-9, 6-11
- Oracle Messaging Gateway
 - checking requirements on Tru64 UNIX, 4-7
 - CSD requirements
 - on AIX, 2-8
 - on HP-UX, 3-8
 - on Linux, 5-8
 - on Solaris, 6-10
 - on Tru64 UNIX, 4-8
 - requirements on AIX, 2-6
 - requirements on HP-UX, 3-6
 - requirements on Linux, 5-6
 - requirements on Solaris, 6-7
 - requirements on Tru64 UNIX, 4-6
- Oracle Net
 - identifying Oracle home for listener, 2-40, 3-37, 4-31, 5-43, 6-41
 - lsnrctl command, 2-41, 3-37, 4-32, 5-43, 6-41
 - stopping existing listener, 2-40, 3-36, 4-31, 5-43, 6-41
 - stopping listener, 2-40, 3-36, 4-31, 5-43, 6-41
 - stopping the listener, 2-41, 3-37, 4-32, 5-43, 6-41
- Oracle Software Owner user
 - and ASM disks, 2-30, 3-27, 4-30, 5-35, 5-36, 5-40, 6-34
 - configuring environment for, 2-41, 3-37, 4-32, 5-44, 6-42
 - creating, 2-12, 3-12, 4-11, 5-11, 6-12, 6-13
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-10, 3-10, 4-9, 5-9, 6-11
 - determining default shell, 2-41, 3-38, 4-33, 5-44, 6-42
 - raw device owner
 - on AIX, 2-36, 2-37
 - on HP-UX, 3-33

- on Linux, 5-39
 - on Solaris, 6-38
- relationship with Oracle base directory, 2-17, 3-18, 4-17, 5-19, 6-20
- required group membership, 2-10, 3-10, 4-9, 5-9, 6-11
- setting shell limits for on Linux, 5-18
- setting up user equivalence, 2-15, 3-15, 4-14, 6-17
- Oracle Spatial
 - Motif requirement on AIX, 2-6
 - Motif requirement on HP-UX, 3-6
 - Motif requirement on Solaris, 6-7
 - Motif requirement on Tru64 UNIX, 4-6
 - requirement for example programs
 - on AIX, 2-6
 - on HP-UX, 3-6
 - on Linux, 5-6
 - on Solaris, 6-7
 - on Tru64 UNIX, 4-6
 - X window requirement on Linux, 5-6
- Oracle Transparent Gateway
 - Informix requirements on HP-UX, 3-7
 - Informix requirements on Solaris, 6-8
 - Ingres requirements on HP-UX, 3-7
 - Ingres requirements on Solaris, 6-8
 - Sybase requirements on AIX, 2-7
 - Sybase requirements on HP-UX, 3-7
 - Sybase requirements on Solaris, 6-7
 - Sybase requirements on Tru64 UNIX, 4-6
 - Teradata requirements on HP-UX, 3-7
 - Teradata requirements on Solaris, 6-8
- Oracle Universal Installer
 - overview of processing, 1-4
- oracle user
 - and ASM disks, 2-30, 3-27, 4-30, 5-35, 5-36, 5-40, 6-34
 - configuring environment for, 2-41, 3-37, 4-32, 5-44, 6-42
 - creating, 2-12, 3-12, 4-11, 5-11, 6-12, 6-13
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-10, 3-10, 4-9, 5-9, 6-11
 - determining default shell, 2-41, 3-38, 4-33, 5-44, 6-42
 - raw device owner
 - on AIX, 2-36, 2-37
 - on HP-UX, 3-33
 - on Linux, 5-39
 - relationship with Oracle base directory, 2-17, 3-18, 4-17, 5-19, 6-20
 - required group membership, 2-10, 3-10, 4-9, 5-9, 6-11
 - setting shell limits for on Linux, 5-18
 - setting up user equivalence, 2-15, 3-15, 4-14, 6-17
- ORACLE_BASE environment variable, 2-18, 2-21, 3-18, 3-22, 4-17, 4-21, 5-19, 5-23, 6-20, 6-24
- removing from shell startup file, 2-42, 3-38, 4-33, 5-44, 6-42
- setting, 2-41, 3-37, 4-32, 5-44, 6-42
- ORACLE_HOME environment variable
 - removing from shell startup file, 2-42, 3-38, 4-33, 5-44, 6-42
 - setting, 2-41, 3-37, 4-32, 5-43, 6-41
 - unset, 2-43, 3-39, 4-34, 5-46, 6-44
- ORACLE_SID environment variable
 - removing from shell startup file, 2-42, 3-38, 4-33, 5-44, 6-42
 - setting, 2-41, 3-37, 4-32, 5-44, 6-42
- oraInst.loc file, 2-19, 3-20, 4-19, 5-21, 6-22
 - location, 5-10
- oraInventory directory
 - See Oracle Inventory directory
- oratab file
 - format, 2-19, 3-20, 4-19, 5-21, 6-22
 - location of, 2-19, 3-20, 4-19, 5-21, 6-22
- OSBDA group
 - and ASM disks on AIX, 2-30
 - and ASM disks on HP-UX, 3-27
 - and ASM disks on Linux, 5-35, 5-36, 5-40
 - and ASM disks on Solaris, 6-34
 - and ASM disks on Tru64 UNIX, 4-30
- OSDBA group
 - and SYSDBA privilege, 2-9, 3-10, 4-8, 5-9, 6-10
 - creating, 2-11, 3-11, 4-10, 5-11, 6-12
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-9, 3-10, 4-8, 5-9, 6-10
 - granting privileges on HP-UX, 3-15
 - raw device group
 - on AIX, 2-36, 2-37
 - on HP-UX, 3-33
 - on Linux, 5-39
 - on Solaris, 6-38
- oslevel command, 2-7
- OSOPER group
 - and SYSOPER privilege, 2-9, 3-10, 4-8, 5-9, 6-10
 - creating, 2-11, 3-12, 4-10, 5-11, 6-12
 - creating on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 - description, 2-9, 3-10, 4-8, 5-9, 6-10

P

- packages
 - checking on Linux, 5-7
 - checking on Solaris, 6-8
 - requirements for Solaris, 6-7
- parameters
 - initialization, 11-1
- partition
 - using with ASM, 1-5, 2-27, 3-24, 4-26, 5-28, 6-30
- partitions
 - binding to raw devices on Linux, 5-39
 - creating on Linux, 5-38
 - creating raw partitions on Linux, 5-37
 - identifying on Linux, 5-37
 - required sizes for raw devices on Linux, 5-38
 - slices on Solaris, 6-34, 6-35
- passwd command, 2-13, 2-14, 3-13, 3-14, 4-12, 4-13, 5-12, 5-14, 6-13, 6-15

- password file
 - raw device for
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- patchadd command, 6-10
- patches
 - checking on AIX, 2-8
 - checking on HP-UX, 3-8
 - checking on Linux, 5-8
 - checking on Solaris, 6-10
 - download location for AIX, 2-8
 - download location for HP-UX, 3-9
 - download location for Linux, 5-8
 - download location for Solaris, 6-10
 - download location for Tru64 UNIX, 4-8
 - requirements on AIX, 2-7
 - requirements on HP-UX, 3-7
 - requirements on Linux, 5-8
 - requirements on Solaris, 6-8
 - requirements on Tru64 UNIX, 4-7
- PATH environment variable
 - description and setting, 2-43, 3-39, 5-46, 6-44
 - setting, 2-41, 3-37, 4-32, 5-44, 6-42
- PC X server
 - installing from, 2-2, 3-2, 4-2, 5-2, 6-2
- per_proc_address_space subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- per_proc_data_size subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- per_proc_stack_size subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- permissions
 - for datafile directories, 2-25, 4-25, 5-27, 6-29
 - for Oracle base directory, 2-21, 3-21, 4-21, 5-22, 6-24
- physical RAM requirements, 2-3, 3-3, 4-3, 5-3, 6-3
- pkginfo command, 6-8
- platforms
 - checking architecture support, 2-4, 3-4, 5-4, 6-4
- Precompilers
 - requirements on AIX, 2-7
 - requirements on HP-UX, 3-6
 - requirements on Linux, 5-7
 - requirements on Tru64 UNIX, 4-6
- preconfigured database
 - ASM disk space requirements, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
 - requirements when using ASM, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
- preconfigured database installation types, 8-2
- Pre-connect
 - TAF failover policy, 9-3
- Preferred
 - service configuration policy, 9-2
- pre-installation
 - raw device creation, C-2
- privgroup file, 3-15
- privileged groups
 - for Oracle Database, 2-9, 3-10, 4-8, 5-9, 6-10
- Pro*C/C++
 - patches required on HP-UX, 3-8
 - patches required on Solaris, 6-9
 - PTFs and APARs required on AIX, 2-8
 - requirements on AIX, 2-7
 - requirements on HP-UX, 3-6
 - requirements on Linux, 5-7
 - requirements on Tru64 UNIX, 4-6
- Pro*COBOL
 - environment variables required for, 2-43, 3-39, 5-46, 6-44
- proc
 - kernel subsystem on Tru64 UNIX, 4-15
- process
 - stopping existing, 2-40, 3-36, 4-31, 5-43, 6-41
 - stopping existing listener process, 2-40, 3-36, 4-31, 5-43, 6-41
 - stopping listener process, 2-40, 3-36, 4-31, 5-43, 6-41
- processor
 - checking system architecture, 2-4, 3-4, 5-4, 6-4
 - /proc/sys/fs/file-max file, 5-16
 - /proc/sys/kernel/sem file, 5-16
 - /proc/sys/kernel/shmall file, 5-16
 - /proc/sys/kernel/shmmni file, 5-16
 - /proc/sys/net/ipv4/ip_local_port_range file, 5-16
 - .profile file, 2-41, 3-38, 4-33, 5-44, 6-42
 - profile file, 5-18
 - profile.local file, 5-18
- program technical fix
 - See PTF
- ps command, 2-40, 3-37, 4-31, 5-43, 6-41
- PTF
 - checking, 2-9
 - download location, 2-9
 - requirements on AIX, 2-7
- pvcreate command, 3-31
- pvdisplay command, 3-27, 3-30

Q

- quality pack
 - checking requirements on HP-UX, 3-8
 - download location for HP-UX, 3-8
 - requirements on HP-UX, 3-8

R

- RAC
 - configuring disks for ASM on HP-UX, 3-27
 - configuring disks for ASM on Linux, 5-33, 5-34
 - configuring disks for ASM on Solaris, 6-33
 - configuring disks for ASM on Tru64 UNIX, 4-29
 - configuring disks for raw devices on HP-UX, 3-30
 - configuring disks for raw devices on Linux, 5-37
 - configuring raw logical volumes on AIX, 2-32
 - configuring raw logical volumes on Solaris, 6-35
 - file system requirements, 2-23, 2-24, 4-22, 4-23,

- 4-24, 5-24, 5-25, 6-26, 6-27, 6-28
- patches required on HP-UX, 3-8
- preferred datafile storage mechanisms, 2-26, 3-23, 4-25, 5-27, 6-29
- requirements for Oracle base directory, 2-21, 3-21, 4-20, 5-22, 6-23
- supported storage locations, 2-22, 4-22, 5-24, 6-26
- supported storage options, 2-22, 3-23, 4-22, 5-24, 6-26

RAC high availability extensions

- configuring services for, 9-2
- TAF policies for, 9-3

rad_gh_regions attribute

- recommended value on Tru64 UNIX, 4-15

RAID

- compared to ASM, 1-5
- device names on Linux, 5-33, 5-35, 5-37
- recommended ASM redundancy level, 1-6, 2-27, 3-24, 4-27, 5-28, 6-30
- using for Oracle datafiles, 2-23, 4-23, 5-25, 6-27

RAM requirements, 2-3, 3-3, 4-3, 5-3, 6-3

raw command, 5-35, 5-36, 5-39

raw device

- for OCR
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for password file
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for server parameter file
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for SPFILE
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for voting disk
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37

raw devices

- activating volume group on cluster nodes
 - on AIX, 2-37
- binding on Linux, 5-35, 5-36
- binding to partitions on Linux, 5-39
- block and character device names on
 - HP-UX, 3-30
 - block and character device names on Solaris, 6-35
- changing permissions and owner for ASM on
 - Linux, 5-35, 5-36, 5-40
- checking availability of disks on AIX, 2-32
- checking disk availability on HP-UX, 3-30
- checking disk availability on Solaris, 6-35
- configuring on AIX, 2-31
- configuring on HP-UX, 3-28
- configuring on Solaris, 6-34
- creating a raw device mapping file
 - on AIX, 2-37
 - on HP-UX, 3-34
 - on Linux, 5-40
 - on Solaris, 6-38
- creating logical volumes on AIX, 2-35
- creating logical volumes on HP-UX, 3-32
- creating logical volumes on Solaris, 6-37
- creating partitions on Linux, 5-38
- creating raw logical volumes on AIX, 2-34
- creating raw logical volumes on HP-UX, 3-31
- creating raw logical volumes on Solaris, 6-36
- creating raw partitions on Linux, 5-37
- DBCA_RAW_CONFIG environment
 - variable, 2-43, 3-39, 5-46, 6-44
- device names on Linux, 5-35, 5-36, 5-39
 - device names, 5-39
- displaying attached disks
 - on Linux, 5-37
- for control files
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for EXAMPLE tablespace
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36
- for redo log files
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for SYSAUX tablespace
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36
- for SYSTEM tablespace
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36
- for TEMP tablespace
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- for UNDOTBS tablespace
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36
- for USER tablespace
 - on AIX, 2-34

- on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- identifying available disks on Linux, 5-37
- identifying disks on AIX, 2-32
- identifying disks on HP-UX, 3-30
- identifying disks on Linux, 5-37
- identifying disks on Solaris, 6-35
- identifying LVM disks on HP-UX, 3-30
- importing on disk group on cluster nodes
 - on AIX, 2-36
 - on HP-UX, 3-33
 - on Solaris, 6-38
- initializing disks for LVM on AIX, 2-30, 2-33
- initializing disks for LVM on HP-UX, 3-31
- initializing disks for LVM on Solaris, 6-36
- required sizes
 - on AIX, 2-34
 - on HP-UX, 3-31
 - on Linux, 5-38
 - on Solaris, 6-36
- setting up, C-2
- specifying location of mapping file, 2-43, 3-39, 5-46, 6-44
- specifying owner and permissions on AIX, 2-36, 2-37
- specifying owner and permissions on HP-UX, 3-33
- specifying owner and permissions on Linux, 5-39
- specifying owner and permissions on Solaris, 6-38
- storage option for datafiles, 2-22, 3-22, 4-21, 5-23, 6-24
- storage option for RAC, 2-22, 3-23, 4-22, 5-24, 6-26
- value for DBCA_RAW_CONFIG environment variable
 - on AIX, 2-39
 - on HP-UX, 3-35
 - on Linux, 5-41
 - on Solaris, 6-40
- verification, 9-3
- raw file, 5-39
- raw logical volumes
 - See* raw devices
- raw partitions
 - See* raw devices
- rawdevices file, 5-35, 5-36, 5-39
- RBS tablespace
 - description, 12-3
- rcp command, 2-15, 3-15, 4-14, 6-17
- rdg
 - kernel subsystem on Tru64 UNIX, 4-15
- rdg_auto_msg_wires subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- rdg_max_auto_msg_wires subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- Real Application Clusters
 - components, 1-7
 - databases, deleting, 9-9
 - installation requirements, 1-2
 - installed components of, 1-8
 - overview, 1-1, 9-1, 12-1
 - raw device setup, C-2
 - See* RAC
- reboot command, 6-19
- recovery files
 - options for placing on file system, 2-23, 4-22, 5-24, 6-26
- Red Hat
 - binding partitions to raw devices, 5-39
 - checking glibc errata, 5-8
 - kernel errata download location, 5-8
 - operating system requirements, 5-6
 - required errata, 5-8
- Red Hat Package Manager
 - See* RPM
- redo log files, 1-8
 - described, 12-3
 - raw devices for
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- redundancy level
 - and space requirements for preconfigured database, 1-7, 2-27, 3-24, 4-27, 5-29, 6-31
 - for ASM, 1-6
- redundant array of independent disks
 - See* RAID
- registration
 - cross-node, 12-4
- remote shell
 - configuring user equivalence for, 2-15, 3-15, 4-14, 6-17
- remsh command, 2-15, 3-15, 4-14, 6-18
- Rendezvous
 - requirement on HP-UX, 3-6
 - requirement on Solaris, 6-7
- requirements
 - hardware, 2-3, 3-3, 4-3, 5-3, 6-3
- .rhosts file, 2-15, 3-15, 4-14, 6-17
- rollback segments
 - described, 12-3
- root user
 - logging in as, 2-2, 3-2, 4-2, 5-2, 6-2
- RPM
 - checking on Linux, 5-7
- rpm command, 5-7, 5-8
- rsh command, 4-14
 - configuring user equivalence for, 2-15, 3-15, 4-14, 6-17
- rt
 - kernel subsystem on Tru64 UNIX, 4-15

S

- SAM
 - starting on HP-UX, 3-17
- sam command, 3-17

- Save as a Database Template
 - option on Creation options page, 9-8
- scp command, 2-15, 3-15, 4-14, 6-17
- SCSI disks
 - device names on Linux, 5-33, 5-35, 5-37
- scstat command, 6-40
- secure shell
 - as alternative to remote shell (rsh), 2-15, 3-15, 4-14, 6-17
- sem file, 5-16
- seminfo_semmni parameter
 - recommended value on Solaris, 6-18
- seminfo_semmns parameter
 - recommended value on Solaris, 6-18
- seminfo_semmsl parameter
 - recommended value on Solaris, 6-18
- seminfo_semvmx parameter
 - recommended value on Solaris, 6-18
- semmap parameter
 - recommended value on HP-UX, 3-17
- semmni parameter
 - recommended value on HP-UX, 3-17
 - recommended value on Linux, 5-16
 - recommended value on Solaris, 6-18
- semmns parameter
 - recommended value on HP-UX, 3-17
 - recommended value on Linux, 5-16
 - recommended value on Solaris, 6-18
- semmnu parameter
 - recommended value on HP-UX, 3-17
- semmsl parameter
 - recommended value on Linux, 5-16
 - recommended value on Solaris, 6-18
- semopm parameter
 - recommended value on Linux, 5-16
- semvmx parameter
 - recommended value on HP-UX, 3-17
 - recommended value on Solaris, 6-18
- server parameter file
 - raw device for
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- server parameter files, 1-8, 11-1, 11-2, 12-4
 - errors, 11-4
- service command, 5-36, 5-39
- Service Management
 - using, 9-2
- service registration
 - configuring, 12-4
- Serviceguard
 - checking on HP-UX, 3-36
 - HP-UX requirement, 3-6
 - starting a cluster node, 3-36
 - starting on HP-UX, 3-36
- services, 9-2
- setld command, 4-7
- setprivgrp command, 3-16
- shared configuration file, 12-2
- shared server, 12-4
- shared storage
 - requirement for RAC, 2-22, 4-22, 5-24, 6-26
- shell
 - default shell startup file, 5-18
 - determining default shell for oracle user, 2-41, 3-38, 4-33, 5-44, 6-42
- SHELL environment variable
 - checking value of, 2-41, 3-38, 4-33, 5-44, 6-42
- shell limits
 - setting on Linux, 5-18
- shell startup file
 - editing, 2-41, 3-38, 4-33, 5-44, 6-42
 - removing environment variables, 2-42, 3-38, 4-33, 5-44, 6-42
- shm_max subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- shm_min subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- shm_mni subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- shm_seg subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- shmall file, 5-16
- shmall parameter
 - recommended value on Linux, 5-16
- shminfo_shmmax parameter
 - recommended value on Solaris, 6-18
- shminfo_shmmin parameter
 - recommended value on Solaris, 6-18
- shminfo_shmmni parameter
 - recommended value on Solaris, 6-18
- shminfo_shmseg parameter
 - recommended value on Solaris, 6-18
- shmmax file, 5-16
- shmmax parameter
 - recommended value on HP-UX, 3-17
 - recommended value on Linux, 5-16
 - recommended value on Solaris, 6-18
- shmmin parameter
 - recommended value on Solaris, 6-18
- shmmni file, 5-16
- shmmni parameter
 - recommended value on HP-UX, 3-17
 - recommended value on Linux, 5-16
 - recommended value on Solaris, 6-18
- shmseg parameter
 - recommended value on HP-UX, 3-17
 - recommended value on Solaris, 6-18
- shutdown command, 4-16
- SID
 - setting ORACLE_SID environment variable, 2-41, 3-37, 4-32, 5-44, 6-42
- sid Prefix, 8-6
- sizer command, 4-7
- smit command, 2-11, 2-12, 2-14
- Software Only
 - configuration type, 8-2
- software requirements
 - AIX, 2-6

- checking on AIX, 2-7
- checking on HP-UX, 3-7
- checking on Linux, 5-7
- checking on Solaris, 6-8
- checking on Tru64 UNIX, 4-6
 - HP-UX, 3-6
 - Linux, 5-6
 - Solaris, 6-6
 - Tru64 UNIX, 4-5
- Solaris
 - block and character device names, 6-34, 6-35
 - checking disk availability for ASM, 6-33
 - checking disk availability for raw devices, 6-35
 - checking existence of nobody user, 6-14
 - checking kernel parameters, 6-19
 - checking patch requirements, 6-10
 - checking software requirements, 6-8
 - checking Sun Cluster software, 6-40
 - checking system architecture, 6-4
 - checking version, 6-8
 - configuring disks for ASM, 6-33
 - configuring kernel parameters, 6-18
 - configuring raw devices, 6-34
 - creating a raw device mapping file, 6-38
 - creating a volume group, 6-35
 - creating logical volumes, 6-37
 - creating raw logical volumes, 6-36
 - creating volume groups, 6-36
 - font packages for Java, 6-7
 - identifying disks for ASM, 6-33
 - identifying disks for LVM, 6-35
 - identifying file systems, 6-34, 6-35
 - identifying LVM disks, 6-34, 6-35
 - importing raw device disk group, 6-38
 - initializing disks for LVM, 6-36
 - location of oratab file, 6-22
 - making kernel parameter changes persist, 6-19
 - operating system requirements, 6-7
 - package requirements, 6-7
 - patch download location, 6-10
 - patch requirements, 6-8
 - raw device sizes, 6-36
 - software requirements, 6-6
 - specifying owner and permissions on raw devices, 6-38
 - Sun Cluster requirement, 6-7
 - verifying that cluster software is running, 6-40
 - Veritas Volume Manager, 6-34, 6-35
- SPFILE
 - raw device for
 - on AIX, 2-35
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-37
- sqlnet.ora file, 12-11
 - default configuration, 12-11
- SSH
 - as alternative to remote shell (rsh), 2-15, 3-15, 4-14, 6-17
- ssm_threshold subsystem attribute
 - recommended value on Tru64 UNIX, 4-15
- startup file
 - default shell startup file, 5-18
 - for shell, 2-41, 3-38, 4-33, 5-44, 6-42
- storage device
 - shared storage requirements, 2-22, 4-22, 5-24, 6-26
- Storage Options page, 9-4
- subsets
 - checking on Tru64 UNIX, 4-7
 - requirements for Tru64 UNIX, 4-6
- Summary dialog, 9-9
- Sun Cluster
 - checking on Solaris, 6-40
 - patches required on Solaris, 6-9
 - requirement on Solaris, 6-7
- swap space
 - requirements, 2-3, 3-3, 4-3, 5-3, 6-3
- swlist command, 3-7, 3-8, 3-9
- Sybase Adaptive Server Enterprise
 - requirements on AIX, 2-7
 - requirements on HP-UX, 3-7
 - requirements on Solaris, 6-7
 - requirements on Tru64 UNIX, 4-6
- Sybase Open Client
 - requirements on AIX, 2-7
 - requirements on HP-UX, 3-7
 - requirements on Solaris, 6-7
 - requirements on Tru64 UNIX, 4-6
- symbolic links
 - X library links required on HP-UX, 3-9
- SysAUX tablespace
 - raw device for
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36
- sysconfig command, 4-16
- sysconfigdb command, 4-16
- sysconfigtab file, 4-16
- sysctl command, 5-17
- sysctl.conf file, 5-17
- SYSDBA privilege
 - associated UNIX group, 2-9, 3-10, 4-8, 5-9, 6-10
- sysdef command, 6-19
- SYSOPER privilege
 - associated UNIX group, 2-9, 3-10, 4-8, 5-9, 6-10
- System Administration Manager
 - See SAM
- system architecture
 - checking, 2-4, 3-4, 5-4, 6-4
- system file, 6-19
- SYSTEM tablespace
 - description, 12-3
 - raw device for
 - on AIX, 2-34
 - on HP-UX, 3-32
 - on Linux, 5-38
 - on Solaris, 6-36

T

tablespaces
 and the DBCA, 12-2
 expanding for large sorts, 12-3
 RBS, 12-3
 SYSTEM, 12-3
 TEMP, 12-3
 undo tablespaces for automatic undo management, 12-3
 USERS, 12-3

tcsh shell
 setting shell limits on Linux, 5-18

TEMP environment variable, 2-3, 3-3, 4-3, 5-3, 6-3
 setting, 2-42, 3-38, 4-33, 5-45, 6-43

TEMP tablespace
 described, 12-3
 raw device for
 on AIX, 2-35
 on HP-UX, 3-32
 on Linux, 5-38
 on Solaris, 6-37

temporary directory, 2-3, 3-3, 4-3, 5-3, 6-3

temporary disk space
 checking, 2-3, 3-3, 4-3, 5-3, 6-3
 freeing, 2-3, 3-3, 4-3, 5-3, 6-3
 requirements, 2-3, 3-3, 4-3, 5-3, 6-3

Teradata ODBC Driver
 requirements on HP-UX, 3-7
 requirements on Solaris, 6-8

TIBCO Rendezvous
 requirement on HP-UX, 3-6
 requirement on Solaris, 6-7

/tmp directory
 checking space in, 2-3, 3-3, 4-3, 5-3, 6-3
 freeing space in, 2-3, 3-3, 4-3, 5-3, 6-3

TMPDIR environment variable, 2-3, 3-3, 4-3, 5-3, 6-3
 setting, 2-42, 3-38, 4-33, 5-45, 6-43

TNS_ADMIN environment variable
 unsetting, 2-43, 3-39, 4-34, 5-46, 6-44

tnsnames.ora file, 12-7
 default configuration, 12-7

Transaction Processing
 configuration type, 8-2
 description of configuration type, 8-2

Transmission Control Protocol/Internet Protocol (TCP/IP)
 for Real Application Clusters, 1-2

Transparent Application Failover (TAF) policies, 9-2

Transparent Gateway
 See Oracle Transparent Gateway

Tru64 UNIX
 block and character device names, 4-30
 checking disk availability for ASM, 4-29
 checking existence of nobody user, 4-12
 checking kernel subsystem attributes, 4-16
 checking software requirements, 4-6
 checking version, 4-7
 configuring disks for ASM, 4-29
 creating identical users
 UNIX commands
 groupadd, 4-13

 identifying AdvFS file systems, 4-30
 identifying disks for ASM, 4-30
 identifying LSM disks, 4-30
 identifying UFS file systems, 4-30
 identifying unused disks, 4-30
 location of oratab file, 4-19
 making subsystem attribute value changes persist, 4-16
 patch download location, 4-8
 patch requirements, 4-7
 software requirements, 4-5
 subset requirements, 4-6
 TruCluster requirement, 4-6

TruCluster
 requirement for Tru64 UNIX, 4-6

U

UFS
 identifying UFS file systems on Tru64 UNIX, 4-30

uid
 identifying existing, 2-14, 3-13, 4-12, 5-13, 6-14
 specifying, 2-14, 3-14, 4-13, 5-13, 6-14
 specifying on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14

ulimit command, 5-18

umask, 2-43, 3-40, 4-34, 5-46, 6-44

umask command, 2-41, 2-43, 3-37, 3-40, 4-32, 4-34, 5-44, 5-46, 6-42, 6-44

uname command, 3-7, 6-8

undo management, 12-3

UNDOTBS tablespace
 raw device for
 on AIX, 2-34
 on HP-UX, 3-32
 on Linux, 5-38
 on Solaris, 6-36

UnitedLinux
 binding partitions to raw devices, 5-39
 /etc/init.d/boot.local file, 5-17
 operating system requirements, 5-6

UNIX commands, 2-43, 3-40, 4-34, 5-46, 6-44

 cfgmgr, 2-30, 2-32
 chdev, 2-30, 2-33
 chmod, 2-21, 2-25, 2-30, 3-21, 3-27, 4-21, 4-25, 4-30, 5-22, 5-27, 5-35, 5-36, 5-40, 6-24, 6-29, 6-34
 chown, 2-21, 2-25, 2-30, 3-21, 3-27, 4-21, 4-25, 4-30, 5-22, 5-27, 5-35, 5-36, 5-40, 6-24, 6-29, 6-34
 clic_stat, 3-36
 clu_get_info, 4-31
 cmruncl, 3-36
 cmrunnode, 3-36
 cmviewcl, 3-36
 disklabel, 4-30
 dupatch, 4-8
 env, 2-43, 3-40, 4-34, 5-46, 6-44
 fdisk, 5-33, 5-35, 5-37, 5-38
 file, 4-29
 format, 6-33, 6-35
 getconf, 2-4, 3-4, 5-4, 6-4

getprivgrp, 3-16
 groupadd, 5-11
 id, 2-13, 2-14, 3-13, 4-12, 5-12, 5-13, 6-14
 importvg, 2-36
 insmod, 5-42
 instfix, 2-8
 ioscan, 3-27, 3-30
 isainfo, 2-4, 3-4, 5-4, 6-4
 limit, 5-18
 load_ocfs, 5-42
 lsdev, 2-30, 2-32, 5-33, 5-35
 lslpp, 2-7, 2-9
 lsmod, 5-42
 lspv, 2-30, 2-32, 2-36
 lvcreate, 3-32
 mkdir, 2-21, 2-25, 3-21, 4-21, 4-25, 5-22, 5-27, 6-24, 6-29
 mklv, 2-35
 mknod, 3-31
 mkvg, 2-33
 mount, 4-30
 oslevel, 2-7
 passwd, 2-13, 2-14, 3-13, 3-14, 4-12, 4-13, 5-12, 5-14, 6-13, 6-15
 patchadd, 6-10
 pkginfo, 6-8
 ps, 2-40, 3-37, 4-31, 5-43, 6-41
 pvcreate, 3-31
 pvdisplay, 3-27, 3-30
 raw, 5-35, 5-36, 5-39
 rcp, 2-15, 3-15, 4-14, 6-17
 reboot, 6-19
 remsh, 2-15, 3-15, 4-14, 6-18
 rpm, 5-7, 5-8
 rsh, 4-14
 sam, 3-17
 scp, 2-15, 3-15, 4-14, 6-17
 scstat, 6-40
 service, 5-36, 5-39
 setid, 4-7
 setprivgrp, 3-16
 shutdown, 4-16
 sizer, 4-7
 smit, 2-11, 2-12, 2-14
 swap, 2-3, 3-3, 4-3, 5-3, 6-3
 swapon, 2-3, 3-3, 4-3, 5-3, 6-3
 swlist, 3-7, 3-8, 3-9
 sysconfig, 4-16
 sysconfigdb, 4-16
 sysctl, 5-17
 sysdef, 6-19
 ulimit, 5-18
 umask, 2-41, 3-37, 4-32, 5-44, 6-42
 uname, 3-7, 6-8
 unset, 2-43, 3-39, 4-34, 5-46, 6-44
 useradd, 3-13, 3-14, 4-11, 4-13, 5-12, 5-13, 6-13, 6-15
 varyoffvg, 2-36
 varyonvg, 2-34
 vgchange, 3-31, 3-33
 vgcreate, 3-31
 vgexport, 3-33
 vgmport, 3-33
 voldisk, 4-30
 vxassist, 6-37
 vxdctl, 6-38
 vxdg, 6-36, 6-38
 vxdisk, 6-34, 6-35
 vxdiskadd, 6-36
 vxdiskconfig, 6-35
 vxvol, 6-38
 xhost, 2-2, 3-2, 4-2, 5-2, 6-2
 xterm, 2-2, 3-2, 4-2, 5-2, 6-2

UNIX groups
 checking for existing oinstall group, 2-10, 3-11, 4-10, 5-10, 6-12
 creating identical groups on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 creating the dba group, 2-11, 3-11, 4-10, 5-11, 6-12
 creating the oinstall group, 2-10, 3-11, 4-10, 5-10, 6-11
 creating the oper group, 2-11, 3-12, 4-10, 5-11, 6-12
 oinstall, 2-10, 3-10, 4-9, 5-9, 6-11
 OSDBA (dba), 2-9, 3-10, 4-8, 5-9, 6-10
 OSOPER (oper), 2-9, 3-10, 4-8, 5-9, 6-10
 required for oracle user, 2-10, 3-10, 4-9, 5-9, 6-11
 requirements, 2-9, 3-9, 4-8, 5-9, 6-10
 specifying when creating users, 2-14, 3-14, 4-13, 5-13, 6-14
 using NIS, 2-10, 2-13, 3-11, 3-13, 4-9, 4-12, 5-10, 5-13, 6-11, 6-14

UNIX users
 checking existence of the nobody user, 2-13, 4-12, 5-12, 6-14
 creating an unprivileged user on HP-UX, 3-13
 creating identical users on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 creating the oracle user, 2-12, 3-12, 4-11, 5-11, 6-12, 6-13
 nobody, 2-9, 4-9, 5-9, 6-10
 oracle, 2-10, 3-10, 4-9, 5-9, 6-11
 required for external jobs, 2-9, 3-10, 4-9, 5-9, 6-10
 requirements, 2-9, 3-9, 4-8, 5-9, 6-10
 setting shell limits for on Linux, 5-18
 setting up user equivalence, 2-15, 3-15, 4-14, 6-17
 specifying groups when creating, 2-14, 3-14, 4-13, 5-13, 6-14
 unprivileged user, 2-9, 3-10, 4-9, 5-9, 6-10
 using NIS, 2-10, 2-13, 3-11, 3-13, 4-9, 4-12, 5-10, 5-13, 6-11, 6-14

UNIX workstation
 installing from, 2-2, 3-2, 4-2, 5-2, 6-2

unprivileged user
 checking on AIX, 2-13
 checking on Linux, 5-12
 checking on Solaris, 6-14
 checking Tru64 UNIX, 4-12
 creating on HP-UX, 3-13

unset command, 2-43, 3-39, 4-34, 5-46, 6-44

User Datagram Protocol (UDP)
 for Real Application Clusters, 1-2

user equivalence
 /etc/hosts.equiv file, 2-15, 3-15, 4-14, 6-17
 .rhosts file, 2-15, 3-15, 4-14, 6-17
 setting up, 2-15, 3-15, 4-14, 6-17
 testing, 2-15, 3-15, 4-14, 6-18

user IDs
 identifying existing, 2-14, 3-13, 4-12, 5-13, 6-14
 specifying, 2-14, 3-14, 4-13, 5-13, 6-14
 specifying on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14

USER tablespace
 raw device for
 on AIX, 2-34
 on HP-UX, 3-32
 on Linux, 5-38
 on Solaris, 6-37

useradd command, 3-13, 3-14, 4-11, 4-13, 5-12, 5-13, 6-13, 6-15

users
 checking existence of the nobody user, 2-13, 4-12, 5-12, 6-14
 creating an unprivileged user on HP-UX, 3-13
 creating identical users on other nodes, 2-13, 3-13, 4-12, 5-13, 6-14
 creating the oracle user, 2-12, 3-12, 4-11, 5-11, 6-12, 6-13
 Oracle Software Owner user (oracle), 2-10, 3-10, 4-9, 5-9, 6-11
 setting shell limits for UNIX users on Linux, 5-18
 setting up UNIX user equivalence, 2-15, 3-15, 4-14, 6-17
 specifying UNIX groups when creating, 2-14, 3-14, 4-13, 5-13, 6-14
 UNIX nobody user, 2-9, 4-9, 5-9, 6-10

USERS tablespace
 described, 12-3

V

varyoffvg command, 2-36

varyonvg command, 2-34

Veritas Volume Manager, 6-34, 6-35
 during installation, 6-25

vfs
 kernel subsystem on Tru64 UNIX, 4-15

vgchange command, 3-31, 3-33

vgcreate command, 3-31

vgexport command, 3-33

vgimport command, 3-33

vm
 kernel subsystem on Tru64 UNIX, 4-15

voldisk command, 4-30

volume group
 creating on AIX, 2-32
 creating on HP-UX, 3-30
 creating on Solaris, 6-35
 disabling activation on HP-UX, 3-29

volume groups

 creating on AIX, 2-33
 creating on HP-UX, 3-31
 creating on Solaris, 6-36

volumes
 creating logical volumes on AIX, 2-35
 creating logical volumes on HP-UX, 3-32
 creating logical volumes on Solaris, 6-37
 Veritas Volume Manager, use of during installation, 6-25

voting disk
 raw device for
 on AIX, 2-35
 on HP-UX, 3-32
 on Linux, 5-38
 on Solaris, 6-37
 restriction on using ASM for, 2-26, 3-23, 4-25, 5-27, 6-29

vps_ceiling parameter
 recommended value on HP-UX, 3-17

vxassist command, 6-37

vxctl command, 6-38

vxchg command, 6-36, 6-38

vxdisk command, 6-34, 6-35

vxdiskadd command, 6-36

vxdiskconfig command, 6-35

VxVM, 6-34, 6-35
 creating raw logical volumes on Solaris, 6-36

vxvol command, 6-38

W

WebSphere MQ
 checking requirements on Tru64 UNIX, 4-7

CSD download location
 for AIX, 2-9
 for HP-UX, 3-9
 for Linux, 5-8
 for Solaris, 6-10
 for Tru64 UNIX, 4-8

CSDs required
 on AIX, 2-8
 on HP-UX, 3-8
 on Linux, 5-8
 on Solaris, 6-10
 on Tru64 UNIX, 4-8

requirement on AIX, 2-6
requirement on HP-UX, 3-6
requirement on Linux, 5-6
requirement on Solaris, 6-7
requirement on Tru64 UNIX, 4-6

Welcome page, 9-3

Windows
 analogy for Oracle base directory, 2-17, 3-18, 4-17, 5-19, 6-19

X

X emulator
 installing from, 2-2, 3-2, 4-2, 5-2, 6-2

X library symbolic links

- required on HP-UX, 3-9
- X Window and X/Motif Software development
 - Motif requirement, 4-6
- X window system
 - enabling remote hosts, 2-2, 3-2, 4-2, 5-2, 6-2
- X Windows Motif
 - requirement on Solaris, 6-7
- XFree86-devel
 - requirement on Linux, 5-6
- xhost command, 2-2, 3-2, 4-2, 5-2, 6-2
- xterm command, 2-2, 3-2, 4-2, 5-2, 6-2

