

Oracle® Real Application Clusters

Installation and Configuration Guide

10g Release 1 (10.1) for AIX-Based Systems, Apple Mac OS X,
hp HP-UX, hp Tru64 UNIX, Linux, Solaris Operating System,
and Windows Platforms

Part No. B10766-08

December 2004

Oracle Real Application Clusters Installation and Configuration Guide 10g Release 1 (10.1) for AIX-Based Systems, Apple Mac OS X, hp HP-UX, hp Tru64 UNIX, Linux, Solaris Operating System, and Windows Platforms

Part No. B10766-08

Copyright © 2004, Oracle. All rights reserved.

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

Contributing Authors: Jonathan Creighton, Pat Huey, Raj Kumar

Contributors: Chris Allison, Karin Brandauer, Sudip Datta, Rajiv Jayaraman, Roland Knapp, Diana Lorentz, Barb Lundhild, Vijay Lunawat, John Patrick McHugh, Randy Neville, Michael Polaski, Sudheendra Sampath, Janelle Simmons, Clive Simpkins, Khethavath P. Singh, Nitin Vengurlekar, Gary Young

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.....	xv
Documentation Accessibility	xv
Structure	xvi
Related Documents	xviii
Conventions	xix
What's New in Oracle Database 10g RAC Installation and Configuration?	xxv
Oracle Database 10g Release 1 (10.1) New Features for RAC Installation and Configuration	xxv
Part I Oracle Database 10g Real Application Clusters Installation Planning and Requirements	
1 Introduction to Installing and Configuring Oracle Database 10g RAC	
Real Application Clusters Documentation Overview	1-1
Oracle Real Application Clusters Administrator's Guide.....	1-2
Oracle Real Application Clusters Deployment and Performance Guide	1-2
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	1-3
Cluster Setup and Pre-Installation Configuration Tasks for Real Application Clusters	1-4
Pre-Installation, Installation, and Post-Installation Overview	1-4
Pre-Installation Overview for Oracle Database 10g Real Application Clusters	1-5
Installation Overview for Oracle Database 10g Real Application Clusters.....	1-5
Post-Installation Overview for Oracle Database 10g Real Application Clusters	1-5
The Oracle Universal Installer and Real Application Clusters	1-5
Storage Considerations for Installing Oracle Database 10g Real Application Clusters	1-6
Overview of Automatic Storage Management	1-6
Additional Considerations for Using Oracle Database 10g Features in RAC	1-8
Oracle Database 10g and Real Application Clusters Components	1-9
The Cluster Ready Services Clusterware.....	1-9
The Installed Real Application Clusters Components.....	1-9
Oracle Database 10g Real Application Clusters Version Compatibility	1-10

Required UNIX Groups	1-10
----------------------------	------

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

2 Pre-Installation Tasks for RAC on AIX

Log In to the System as root	2-2
Check the Hardware Requirements.....	2-3
Check the Network Requirements.....	2-4
Check the Software Requirements	2-6
Checking the Software Requirements.....	2-6
Create Required UNIX Groups and User	2-9
Creating the Oracle Inventory Group	2-10
Creating the OSDBA Group	2-11
Creating an OSOPER Group (Optional)	2-11
Creating the Oracle Software Owner User	2-12
Verifying that the UNIX User nobody Exists.....	2-13
Create Identical Users and Groups on Other Cluster Nodes	2-13
Set Up User Equivalence for rsh and rcp on All Cluster Nodes	2-15
Configure Kernel Parameters and Shell Limits	2-15
Configuring Shell Limits, System Configuration, and Network Tuning Parameters.....	2-16
Identify Required Software Directories	2-18
Identify or Create an Oracle Base Directory	2-20
Create the CRS Home Directory.....	2-22
Choose a Storage Option for Oracle CRS, Database, and Recovery Files.....	2-23
Create Directories for Oracle CRS, Database, or Recovery Files	2-24
Configure Disks for Automatic Storage Management	2-28
Identifying Storage Requirements for ASM.....	2-28
Using an Existing ASM Disk Group.....	2-30
Configuring Disks for ASM.....	2-31
Configure Raw Devices.....	2-34
Configuring Raw Disk Devices or Raw Logical Volumes	2-34
Verify the Cluster Software Configuration	2-47
Stop Existing Oracle Processes	2-49
Configure the oracle User's Environment	2-50

3 Pre-Installation Tasks for RAC on HP-UX

Log In to the System as root	3-1
Check the Hardware Requirements.....	3-2
Check the Network Requirements.....	3-3
Check the Software Requirements	3-5
Checking the Software Requirements.....	3-5
Create Required UNIX Groups and User	3-10
Creating the Oracle Inventory Group	3-11
Creating the OSDBA Group	3-12
Creating an OSOPER Group (Optional)	3-12
Creating the Oracle Software Owner User	3-13

Creating an Unprivileged User	3-14
Create Identical Users and Groups on 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	3-16
Configure Kernel Parameters and Shell Limits	3-16
Configuring Kernel Parameters	3-17
Identify Required Software Directories	3-19
Identify or Create an Oracle Base Directory	3-20
Create the CRS Home Directory	3-22
Choose a Storage Option for Oracle CRS, Database, and Recovery Files.....	3-23
Configure Disks for Automatic Storage Management	3-24
Identifying Storage Requirements for ASM.....	3-25
Using an Existing ASM Disk Group.....	3-27
Configuring Disks for ASM.....	3-28
Configure Raw Devices.....	3-29
Configuring Raw Disk Devices or Raw Logical Volumes	3-30
Verify the Cluster Software Configuration.....	3-42
Stop Existing Oracle Processes	3-43
Configure the oracle User's Environment	3-44

4 Pre-Installation Tasks for RAC on hp Tru64 UNIX

Log In to the System as root	4-1
Check the Hardware Requirements.....	4-2
Check the Network Requirements.....	4-3
Check the Software Requirements	4-5
Checking the Software Requirements.....	4-5
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)	4-11
Creating the Oracle Software Owner User	4-11
Verifying that the UNIX User nobody Exists.....	4-12
Create Identical Users and Groups on Other Cluster Nodes	4-12
Set Up User Equivalence for rsh and rcp on All Cluster Nodes	4-14
Configure Kernel Subsystem Attributes	4-15
Configuring Kernel Subsystem Attributes	4-15
Identify Required Software Directories	4-17
Identify or Create an Oracle Base Directory	4-19
Create the CRS Home Directory.....	4-21
Choose a Storage Option for Oracle CRS, Database, and Recovery Files.....	4-22
Create Directories for Oracle CRS, Database, or Recovery Files	4-23
Configure Disks for Automatic Storage Management	4-26
Identifying Storage Requirements for ASM.....	4-27
Using an Existing ASM Disk Group.....	4-29
Configuring Disks for ASM	4-30
Verify that the Cluster Software is Running.....	4-32
Stop Existing Oracle Processes	4-32

Configure the oracle User's Environment	4-33
---	------

5 Pre-Installation Tasks for RAC on Linux

Log In to the System as root	5-2
Check the Hardware Requirements.....	5-3
Check the Network Requirements.....	5-4
Check the Software Requirements	5-5
Checking the Software Requirements on Linux.....	5-6
Create Required UNIX Groups and User	5-11
Creating the Oracle Inventory Group	5-12
Creating the OSDBA Group	5-13
Creating an OSOPER Group (Optional)	5-13
Creating the Oracle Software Owner User.....	5-13
Verifying that the UNIX User nobody Exists.....	5-14
Create Identical Users and Groups on Other Cluster Nodes	5-15
Configure SSH on All Cluster Nodes.....	5-16
Configure Kernel Parameters and Shell Limits	5-18
Configuring Kernel Parameters on Linux	5-18
Identify Required Software Directories	5-21
Identify or Create an Oracle Base Directory	5-23
Create the CRS Home Directory.....	5-26
Choose a Storage Option for Oracle CRS, Database, and Recovery Files.....	5-26
Create Directories for Oracle CRS, Database, or Recovery Files	5-28
Configure Disks for Automatic Storage Management	5-31
Identifying Storage Requirements for ASM.....	5-31
Using an Existing ASM Disk Group.....	5-34
Configuring Disks for ASM.....	5-35
Configuring Disks for ASM on Linux.....	5-35
Configuring Disks for ASM Using the ASM Library Driver.....	5-36
Configuring Disks for ASM Using Raw Devices	5-40
Configure Raw Devices.....	5-44
Configuring Raw Partitions or Raw Logical Volumes on Linux	5-44
Verify that the Required Software is Running.....	5-53
Stop Existing Oracle Processes	5-55
Configure the oracle User's Environment	5-56

6 Pre-installation Tasks for RAC on Mac OS X

Log In to the System as root	6-1
Check the Hardware Requirements.....	6-2
Check the Network Requirements.....	6-3
Check the Software Requirements	6-5
Create Required UNIX Groups and User	6-6
Creating the Oracle Inventory Group	6-7
Creating the OSDBA Group	6-8
Creating an OSOPER Group (Optional)	6-9
Creating the Oracle Software Owner User.....	6-9
Verifying That the UNIX User nobody Exists.....	6-11

Create Identical Users and Groups on Other Cluster Nodes	6-12
Configure SSH on All Cluster Nodes.....	6-14
Configure Kernel Parameters and Shell Limits	6-16
Identify Required Software Directories	6-18
Identify or Create an Oracle Base Directory	6-20
Create the CRS Home Directory.....	6-22
Choose a Storage Option for Oracle CRS, Database, and Recovery Files.....	6-23
Configure Disks for Automatic Storage Management	6-24
Identifying Storage Requirements for ASM.....	6-24
Using an Existing ASM Disk Group.....	6-27
Configuring Disks for ASM.....	6-28
Configure Raw Partitions	6-29
Stop Existing Oracle Processes	6-34
Configure the oracle User's Environment	6-35

7 Pre-Installation Tasks for RAC on Solaris

Log In to the System as root	7-2
Check the Hardware Requirements.....	7-3
Check the Network Requirements.....	7-4
Check the Software Requirements	7-6
Checking the Software Requirements.....	7-6
Create Required UNIX Groups and User	7-11
Creating the Oracle Inventory Group	7-12
Creating the OSDBA Group	7-13
Creating an OSOPER Group (Optional)	7-13
Creating the Oracle Software Owner User	7-13
Verifying that the UNIX User nobody Exists.....	7-14
Create Identical Users and Groups on Other Cluster Nodes	7-15
Set Up User Equivalence for rsh and rcp on All Cluster Nodes	7-16
Configure Kernel Parameters and Shell Limits	7-17
Configuring Kernel Parameters	7-17
Identify Required Software Directories	7-18
Identify or Create an Oracle Base Directory	7-20
Create the CRS Home Directory.....	7-23
Choose a Storage Option for Oracle CRS, Database, and Recovery Files.....	7-23
Create Directories for Oracle CRS, Database, or Recovery Files	7-25
Configure Disks for Automatic Storage Management	7-28
Identifying Storage Requirements for ASM.....	7-28
Using an Existing ASM Disk Group.....	7-31
Configuring Disks for ASM.....	7-32
Configure Raw Partitions or Raw Logical Volumes.....	7-34
Configuring Raw Partitions or Raw Logical Volumes	7-34
Verify that the Cluster Software is Configured and Running.....	7-44
Stop Existing Oracle Processes	7-45
Configure the oracle User's Environment	7-46

8 Pre-Installation Tasks for RAC on Windows

Oracle Database System Requirements	8-1
Software Requirements	8-1
Hardware Requirements	8-3
Hard Disk Space Requirements	8-3
Oracle Cluster File System Pre-Installation Steps	8-4
Verifying Hardware Requirements	8-5
Hardware and Software Certification	8-5
Web Browser Support	8-5
Telnet and Terminal Services Support	8-6
Windows Telnet Services Support	8-6
Windows Terminal Services and Remote Desktop Support	8-6
Network Requirements	8-7
Network Hardware Requirements	8-7
IP Address Requirements	8-7
Checking the Network Requirements	8-8
Individual Component Requirements	8-9
Configuring Disk Storage for Oracle Datafiles and Recovery Files	8-9
Creating Directories for Oracle Datafiles or Recovery Files	8-10
Configuring Disks for Automatic Storage Management	8-12
Identifying Storage Requirements for ASM	8-13
Using an Existing ASM Disk Group	8-15
Configuring Disks for ASM	8-16
Configuring Raw Partitions	8-18
Stop Existing Oracle Services	8-21
Oracle Advanced Security Requirements	8-21
Oracle Enterprise Manager Requirements	8-21
Oracle Managed Files Requirements	8-22
Oracle Transparent Gateway Requirements	8-22

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

9 Installing Cluster Ready Services on UNIX

Installation Setup Procedures	9-1
Installing the UDLM for Sun Clusters	9-2
Installing Cluster Ready Services with the OUI	9-2
Cluster Ready Services Background Processes	9-6

10 Installing Cluster Ready Services on Windows

Prepare to Install Cluster Ready Services on Windows-Based Systems	10-1
Verify Cluster Privileges	10-1
Stop GSD Services from Earlier Releases	10-1
Storage Configuration Steps for Real Application Clusters	10-2
Using Oracle Cluster File System	10-2
Using Raw Devices	10-2

Using Raw Devices for Database When Choosing OCFS for Data Storage During CRS Install	10-2
Using the Oracle Universal Installer to Install Cluster Ready Services on Windows.....	10-3
Formatting Drives to Use Oracle Cluster File System after Installing Oracle Database 10g with RAC	10-7
11 Installing Oracle Database 10g with Real Application Clusters	
Selecting a Database Configuration Type	11-1
Configuration Type Descriptions	11-2
General Purpose, Transaction Processing, and Data Warehouse Configuration Types	11-2
Using the Advanced Configuration Type	11-2
Behavior of the OUI, the DBCA, and Other Assistants During Installation	11-3
Installation Setup Procedures	11-3
Installation Setup Procedures for UNIX-Based Systems.....	11-3
Installation Setup Procedures for Windows-Based Systems	11-4
Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer	11-5
Installation on Windows-Based Systems with the Minimum Memory Requirements	11-12
De-Installing Real Application Clusters Software	11-13
De-Installing Oracle Database 10g RAC Software	11-14
De-Installing Cluster Ready Services	11-17
De-Installing Cluster Ready Services from UNIX Environments.....	11-17
De-Installing Cluster Ready Services from Windows Environments.....	11-18
De-Installing Cluster Ready Services from Windows Environments with No Previous Cluster Software Versions	11-18
De-Installing Oracle Cluster Ready Services from Windows Environments, with Clusterware Downgrade to 9.2	11-18
12 Creating RAC Databases with the Database Configuration Assistant	
Using the Database Configuration Assistant in Real Application Clusters.....	12-1
Benefits of Using the Database Configuration Assistant.....	12-2
Real Application Clusters High Availability Services	12-2
Service Configuration and Instance Preferences	12-2
Transparent Application Failover Policies	12-2
Creating the Database after Installation Using the Database Configuration Assistant	12-2
Creating a Real Application Clusters Database with the DBCA	12-3
Deleting a Real Application Clusters Database with the DBCA	12-9
13 Real Application Clusters Post-Installation Procedures	
Required Post-Installation Tasks	13-1
Back Up the Voting Disk after Installation.....	13-1
Download and Install Patches.....	13-2
Configure Oracle Products	13-2
Oracle Real Application Clusters 10g Installed on an OCFS	13-3
Running Oracle9i RAC with Oracle RAC 10g (Linux Systems Only)	13-3
Register COM Applications (Windows Systems Only).....	13-4
Ensure Valid Path Name Exists (Windows Systems Only)	13-4

Recommended Post-Installation Tasks	13-4
Verifying Enterprise Manager Operations	13-4
Recommended Post-Installation Tasks for UNIX	13-5
Back Up the root.sh Script	13-5
Set Up Users Accounts	13-5
Recommended Post-Installation Task for Windows	13-5
Configure <i>iSQL*Plus</i> for Remote Nodes	13-5
Using Oracle9i Language and Definition Files with Oracle Database 10g Release 1 (10.1)	13-6

Part IV Real Application Clusters Environment Configuration

14 Configuring the Server Parameter File in Real Application Clusters Environments

Parameter Files and Real Application Clusters	14-1
Using Server Parameter Files in Real Application Clusters	14-1
Location of The Server Parameter File	14-2
Parameter File Search Order in Real Application Clusters	14-3
Migrating to the Server Parameter File in Real Application Clusters Environments	14-3
Server Parameter File Placement in Real Application Clusters	14-3
Procedures for Migrating to the Server Parameter File	14-3
Server Parameter File Errors in Real Application Clusters	14-4

15 Understanding the Real Application Clusters Installed Configuration

Understanding the Configured Environment in Real Application Clusters	15-1
The Oracle Cluster Registry in Real Application Clusters	15-1
UNIX oratab Configurations for Real Application Clusters	15-2
Database Components Created Using the Database Configuration Assistant	15-2
Tablespaces and Datafiles	15-2
Control Files	15-4
Redo Log Files	15-4
Managing Undo Tablespaces in Real Application Clusters	15-4
Initialization Parameter Files	15-4
Configuring Service Registration-Related Parameters in Real Application Clusters	15-4
Configuring the Listener File (<i>listener.ora</i>)	15-5
Local Listeners	15-5
Multiple Listeners	15-6
How Oracle Uses the Listener (<i>listener.ora</i> File)	15-6
Listener Registration and PMON Discovery	15-6
Directory Server Access (<i>ldap.ora</i> File)	15-7
Net Service Names (<i>tnsnames.ora</i> File)	15-7
Profile (<i>sqlnet.ora</i> File)	15-12

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-1
	Real Application Clusters Installation Error Messages	A-1
	Performing Cluster Diagnostics During Real Application Clusters Installations	A-1
B	Using Scripts to Create Real Application Clusters Databases	
	Creating a Database Using Scripts.....	B-1
C	Configuring Raw Devices for Real Application Clusters	
	Raw Devices Required by the DBCA for Non-CFS Environments.....	C-1
	Planning Your Raw Device Creation Strategy	C-1
D	Converting to Real Application Clusters from Single-Instance Oracle Databases	
	Deciding to Convert.....	D-1
	Prerequisites for Conversion	D-1
	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-2
	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-4
	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-6
	Post-Conversion Steps.....	D-7
E	Directory Structure for Oracle Database 10g Real Application Clusters Environments	
	Understanding the Real Application Clusters Directory Structure	E-1
	UNIX Directory Structures for Real Application Clusters.....	E-1
	Windows 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, Apple Mac OS X, hp HP-UX, hp Tru64 UNIX, Linux, Solaris Operating System, and Windows Platforms

Part No. B10766-08

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
500 Oracle Parkway, Mailstop 40p11
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

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 RAC on AIX"

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

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

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

Chapter 4, "Pre-Installation Tasks for 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 RAC on Linux"

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

Chapter 6, "Pre-installation Tasks for RAC on Mac OS X"

This chapter describes the pre-installation procedures for installing RAC on Mac OS X-based systems.

Chapter 7, "Pre-Installation Tasks for RAC on Solaris"

This chapter describes the pre-installation procedures for installing RAC on Solaris Operating Systems.

Chapter 8, "Pre-Installation Tasks for RAC on Windows"

This chapter describes the pre-installation procedures for installing RAC on Windows-based 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 on UNIX- and Windows-based systems.

Chapter 9, "Installing Cluster Ready Services on UNIX"

This chapter describes how to install Cluster Ready Services on UNIX-based systems.

Chapter 10, "Installing Cluster Ready Services on Windows"

This chapter describes how to install Cluster Ready Services on Windows-based systems.

Chapter 11, "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 12, "Creating RAC Databases with the Database Configuration Assistant"

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

Chapter 13, "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 14, "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 15, "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 both UNIX- and Windows-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*

Operating System-Specific Administrative Guides

- *Oracle Database Administrator's Reference, 10g Release 1 (10.1) for UNIX Systems*
- *Oracle Database Platform Guide for Windows*
- *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 Database 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://www.oracle.com/technology/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://www.oracle.com/technology/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.

Conventions

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

- [Terminology Conventions](#)
- [Conventions in Text](#)
- [Conventions in Code Examples](#)
- [Conventions for Windows Operating Systems](#)

Terminology Conventions

The names for some operating systems have been shortened in this guide, as described in the following table.

Operating System	Abbreviated Name
AIX-Based Systems	AIX
Apple Mac OS X	Mac OS X
hp HP-UX PA-RISC (64-bit) hp HP-UX Itanium	HP-UX Note: Where the information for HP-UX is different on a particular architecture, this is noted in the text.
hp Tru64 UNIX	Tru64 UNIX
Linux x86 Linux Itanium IBM zSeries Based Linux	Linux Note: Where the information for Linux is different on a particular architecture, this is noted in the text.
Microsoft Windows	Windows
Solaris Operating System (SPARC) Solaris Operating System (x86)	Solaris Note: Where the information for Solaris is different on a particular architecture, this is noted in the text.

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>

Conventions for Windows Operating Systems

The following table describes conventions for Windows operating systems and provides examples of their use.

Convention	Meaning	Example
Choose Start >	How to start a program.	To start the Database Configuration Assistant, choose Start > Programs > Oracle - HOME_NAME > Configuration and Migration Tools > Database Configuration Assistant.

Convention	Meaning	Example
File and directory names	File and directory names are not case sensitive. The following special characters are not allowed: left angle bracket (<), right angle bracket (>), colon (:), double quotation marks ("), slash (/), pipe (), and dash (-). The special character backslash (\) is treated as an element separator, even when it appears in quotes. If the file name begins with \\, then Windows assumes it uses the Universal Naming Convention.	c:\winnt\"system32 is the same as C:\WINNT\SYSTEM32
C:\>	Represents the Windows command prompt of the current hard disk drive. The escape character in a command prompt is the caret (^). Your prompt reflects the subdirectory in which you are working. Referred to as the <i>command prompt</i> in this manual.	C:\oracle\oradata>
Special characters	The backslash (\) special character is sometimes required as an escape character for the double quotation mark (") special character at the Windows command prompt. Parentheses and the single quotation mark (') do not require an escape character. Refer to your Windows operating system documentation for more information on escape and special characters.	C:\>exp scott/tiger TABLES=emp QUERY=\"WHERE job='SALESMAN' and sal<1600\" C:\>imp SYSTEM/password FROMUSER=scott TABLES=(emp, dept)
HOME_NAME	Represents the Oracle home name. The home name can be up to 16 alphanumeric characters. The only special character allowed in the home name is the underscore.	C:\> net start OracleHOME_NAME\TNSListener

Convention	Meaning	Example
<p><i>ORACLE_HOME</i> and <i>ORACLE_BASE</i></p>	<p>In releases prior to Oracle8i release 8.1.3, when you installed Oracle components, all subdirectories were located under a top level <i>ORACLE_HOME</i> directory that by default used one of the following names:</p> <ul style="list-style-type: none"> ■ C:\orant for Windows NT ■ C:\orawin98 for Windows 98 <p>This release complies with Optimal Flexible Architecture (OFA) guidelines. All subdirectories are not under a top level <i>ORACLE_HOME</i> directory. There is a top level directory called <i>ORACLE_BASE</i> that by default is C:\oracle. If you install the latest Oracle release on a computer with no other Oracle software installed, then the default setting for the first Oracle home directory is C:\oracle\orann, where <i>nn</i> is the latest release number. The Oracle home directory is located directly under <i>ORACLE_BASE</i>.</p> <p>All directory path examples in this guide follow OFA conventions.</p> <p>Refer to <i>Oracle Database Platform Guide for Windows</i> for additional information about OFA compliances and for information about installing Oracle products in non-OFA compliant directories.</p>	<p>Go to the <i>ORACLE_BASE\ORACLE_HOME\rdms\admin</i> directory.</p>

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- and Windows-based platforms on which RAC operates.
- The Oracle Database 10g with RAC is available on both the Standard Edition and the Enterprise Edition.

See Also: *Oracle Real Application Clusters Quick Installation Guide for Oracle Database Standard Edition for Windows* for a step-by-step instructions on installing RAC on Windows using the Standard Edition of Oracle.

- 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 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.

Note: DBUA does not support a direct upgrade of Oracle Parallel Server version 8.1.7 databases to Oracle Database 10g with RAC. If you are using Oracle8*i* release 8.1.7, then you can manually upgrade to Oracle Database 10g or use the Oracle9*i* release 9.2 DBUA to upgrade from release 8.1.7 to release 9.2, then use the Oracle Database 10g DBUA to upgrade to the current release.

- Cluster Ready Services (CRS) contains the cluster management software required to support Oracle Database 10g RAC databases. CRS also provides high availability components that provide many system management features. The components of CRS interact with 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 should only be used with Oracle9i databases. The CRS installation process stops any existing GSD processes. To start or stop the GSD processes manually, use `srvctl start nodeapps` or `srvctl stop nodeapps` respectively.
- 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. On Windows-based platforms, the OracleCSService, OracleCRService, and OracleEVMService replace the pre-Oracle Database 10g OracleCMService9i.
- Oracle Database 10g provides cluster file system support for Linux and Windows-based platforms.

See Also: <http://www.oracle.com/technology> 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 Database Administrator's Guide* for information about Automatic Storage Management, a new database file management feature
 - 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.
 - In Windows-based environments using raw partitions, you can use a newly-introduced DBCA raw device mapping file to associate database objects with their partition symbolic link names. This removes the pre-Oracle Database 10g DBCA requirement to always prefix raw partition symbolic links with a database name. This enables you to reuse the same raw partition symbolic links for any database name if that partition is not a part of any existing database.

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 and includes the following topics:

- [Real Application Clusters Documentation Overview](#)
- [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 and Real Application Clusters Components](#)
- [Oracle Database 10g Real Application Clusters Version Compatibility](#)
- [Required UNIX Groups](#)

Real Application Clusters Documentation Overview

This section describes the RAC documentation set. The platform-specific Oracle Database 10g CD contains a copy of this book, the *Oracle Real Application Clusters Installation and Configuration Guide* in both HTML and PDF formats. This book contains the pre-installation, installation, and post-installation information for all UNIX- and Windows-based platforms on which RAC operates. If you are installing Oracle Database 10g Standard Edition with RAC on a Windows-based system, then refer to the *Oracle Real Application Clusters Quick Installation Guide for Oracle Database Standard Edition for Windows*. Additional information for this release may be available in the Oracle Database 10g README or Release Notes.

The Server Documentation CD contains the following additional documentation about RAC administration and deployment: the *Oracle Real Application Clusters Administrator's Guide* and the *Oracle Real Application Clusters Deployment and Performance Guide* as described under the following headings:

- [Oracle Real Application Clusters Administrator's Guide](#)
- [Oracle Real Application Clusters Deployment and Performance Guide](#)

Oracle Real Application Clusters Administrator's Guide

The *Oracle Real Application Clusters Administrator's Guide* provides RAC-specific administration information. Some of the topics described in this book include the use of Oracle Enterprise Manager in RAC environments. The book also describes how to administer services and storage and how to use RAC scalability features to add and delete instances and nodes in RAC environments. The book also discusses how to use Recovery Manager (RMAN) and how to perform backup and recovery in RAC.

The *Oracle Real Application Clusters Administrator's Guide* also describes how to use the Server Control (SRVCTL) utility to start and stop the database and instances, manage configuration information, and to delete or move instances and services. You can also use the appendix to resolve various RAC tools messages. A troubleshooting section describes how to interpret the content of RAC-specific log files.

Oracle Real Application Clusters Deployment and Performance Guide

The *Oracle Real Application Clusters Deployment and Performance Guide* highlights the main deployment topics for RAC by briefly describing Cluster Ready Services (CRS), storage, database creation, and services deployment in RAC. Design and deployment topics in this book describe service topologies and workload management in RAC. Specifically, the book describes how the Automatic Workload Repository tracks and reports service levels and how you can use service level thresholds and alerts to improve high availability in your RAC environment. There is also a services deployment example in the appendix of this book that you can use to learn more about how to deploy and manage services in RAC environments.

The Oracle Real Application Clusters Deployment and Performance Guide provides a high-level description of interconnect protocols, as well as information about how to monitor and tune performance in RAC environments using both Oracle Enterprise Manager and using information in the Automated Workload Repository and Oracle performance views. The book also highlights some application-specific deployment techniques for online transaction processing and data warehousing environments.

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 for storing the Cluster Ready Service and database files.
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. However, note that when the Database Configuration Assistant (DBCA)

configures automatic disk backup, it uses a database recovery area which must be shared. If the database files are stored on a cluster file system, the recovery area can also be shared through the cluster file system. If the database files are stored on an Automatic Storage Management (ASM) disk group, then the recovery area can also be shared through ASM. If the database files are stored on raw devices, a shared directory should be configured using NFS.

- 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.

During an Oracle RAC 10g installation, the Installer displays two screens on which you specify the IP addresses of the private interconnect. The private interconnect is used for inter-node communication by both Oracle CRS and RAC.

On the Cluster Configuration screen, you must specify a private node name to be associated with each public node name. The public node name is the host name of each node, specified as either an alias or an IP address, and the private node name is the private interconnect address that will be used by Oracle CRS. In the Private Node Name field, enter the name of the private interconnect if it is available from a network name server or system hosts file, or else enter the private IP address, unique to each node.

The information on the Private Interconnect Enforcement screen is used to determine which private interconnect will be used by RAC database instances. RAC will use all interconnects identified as private in this page. They must all be in an up state, just as if their IP addresses were specified in the initialization parameter, `CLUSTER_INTERCONNECTS`. RAC does not fail over between cluster interconnects; if one is down then the instances using them will not start.

- 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.

Although not a requirement, you should consider providing redundant switches for all interconnects to increase cluster availability.

While installing and using Real Application Clusters software, you should attempt to keep the system clocks on all of your cluster nodes as close as possible to the same time. Oracle strongly recommends using the Network Time Protocol feature of most operating systems for this purpose, with all nodes using the same reference Network Time Protocol server.

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 Cluster Ready Services (CRS) polling. Your interconnect *must* be certified by Oracle for your platform. You should also have a Web browser to enable Oracle Enterprise Manager and to view online documentation.

For Oracle Database 10g requirements, Oracle's clusterware provides equivalent functionality to vendor clusterware while simplifying installation and reducing support complications. However, vendor clusterware may be needed if you use a

non-ethernet interconnect or if you have deployed clusterware-dependent applications on the same cluster where you deploy RAC.

On Solaris systems, RAC databases on the same cluster must all be 64-bit, as in Oracle Database 10g and Oracle9i Database, or all be 32-bit, as in Oracle9i Database and Oracle8i Database. A mix of 32-bit RAC databases and 64-bit RAC databases on the same cluster is not supported.

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. Some platforms support automatic failover to the alternate interconnect. To enable this, you must configure your operating system-provided failover mechanism.
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. If you use RAC on the Oracle Database 10g Standard Edition, then you must use ASM.

Note: If you use ASM, Oracle recommends that you install ASM in a separate home from the CRS home and the Oracle home, particularly if the ASM instance is to manage storage for more than one RAC database. This reduces downtime when upgrading or de-installing different versions of the software. However, you must create the ASM instance manually because the OUI and DBCA do not support a separate install of ASM.

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 covered in [Part II](#) and [Part III](#) of this book.

Pre-Installation Overview 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 Overview 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 9, "Installing Cluster Ready Services on UNIX"](#) and [Chapter 10, "Installing Cluster Ready Services on Windows"](#). 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 11, "Installing Oracle Database 10g with Real Application Clusters"](#). You use the OUI in this phase to install the RAC software. [Chapter 11, "Installing Oracle Database 10g with Real Application Clusters"](#). The procedures in [Chapter 11](#) 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 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 12, "Creating RAC Databases with the Database Configuration Assistant"](#).

Post-Installation Overview 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 13, "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.

Use the Companion CD to install additional Oracle Database 10g software that may improve performance or extend database capabilities, for example, Oracle JVM, Oracle *interMedia* or Oracle Text.

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

The Oracle Universal Installer and Real Application Clusters

The Oracle Universal Installer (OUI) facilitates the installation of Cluster Ready Services (CRS) and Oracle Database 10g software. In most cases, you use the graphical

user interface (GUI) provided by the OUI to install the software. However, you can also use the OUI to complete scripted installations, without using the GUI. See [Appendix B](#) for information about scripted installations.

When the OUI installs the Oracle software, Oracle recommends that you select a preconfigured database or use the DBCA interactively to create your cluster database. You can also manually create your database as described in procedures posted at <http://www.oracle.com/technology>. 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 create your RAC database using the OUI, or if you do it later using the DBCA, or if you use the Enterprise Manager Configuration Assistant to configure Enterprise Manager, Enterprise Manager is configured for your cluster database. Database control can manage your RAC database, all of its instances, and the hosts where instances are configured.

You can also configure Enterprise Manager Grid Control to manage multiple databases and application servers from a single console. To manage RAC databases in Grid Control, you must install a Grid Control agent on each of the nodes of your cluster. The Agent installation is clusterized, which means you need to perform the install on only one of the cluster nodes.

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.

Note that RAC installations using the Oracle Database Standard Edition must use ASM for database file storage.

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.
- Advanced database

If you choose to create an advanced 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 Components

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 present. If vendor clusterware is not present, then you must use the OUI to enter 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 Mac OS X and 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 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.

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), a component of CRS, on all platforms. The OCSSD performs this function.

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 multiple Oracle homes and different versions of Oracle cluster database software on the same computer as described in the following points:

- You can 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.
- In addition, you cannot install Oracle Database 10g RAC into an existing single-instance Oracle home. If you have an Oracle home for Oracle Database 10g, then use a different Oracle home, and one that is available across the entire cluster for your new installation. Similarly, if you have an Oracle home for an earlier Oracle cluster database software release, then you must also use a different home for the new installation.

If the OUI detects an earlier version of a database, then the OUI asks you about your upgrade preferences. You have the option to upgrade one of the previous-version databases with DBUA or to create a new database using DBCA. The information collected during this dialog is passed to DBUA or DBCA after the software is installed.

- If you run the OUI to install RAC on a system that already has an Oracle Database 10g RAC installation in the same home, then the OUI prompts you to install additional Oracle Database 10g products if you have not already installed all of them.
- You can use the OUI to complete some of the de-install and re-install steps for Oracle Database 10g Real Application Clusters if needed.

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

Required UNIX Groups

This section describes information specific to RAC on UNIX-based platforms. Depending on whether this is the first time you are installing the Oracle server software on your UNIX system, you may need to create several UNIX groups and a UNIX user as described later in the pre-installation procedures. The required UNIX 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 RAC on AIX"](#)
- [Chapter 3, "Pre-Installation Tasks for RAC on HP-UX"](#)
- [Chapter 4, "Pre-Installation Tasks for RAC on hp Tru64 UNIX"](#)
- [Chapter 5, "Pre-Installation Tasks for RAC on Linux"](#)
- [Chapter 6, "Pre-installation Tasks for RAC on Mac OS X"](#)
- [Chapter 7, "Pre-Installation Tasks for RAC on Solaris"](#)
- [Chapter 8, "Pre-Installation Tasks for RAC on Windows"](#)

Pre-Installation Tasks for RAC on AIX

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer (OUI). 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
- Create the CRS Home Directory
- Choose a Storage Option for Oracle CRS, Database, and Recovery Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Configure Raw Devices
- 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-mode 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 silent-mode installations, refer to [Appendix B](#).

- 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 are not installing the software on the local system, use the `ssh`, `rlogin`, or `telnet` commands to connect to the system where you want to install the software:

```
$ 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, refer to 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

Each 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 1 GB or more of memory, do not increase the swap space above 2 GB. Refer to the *Oracle Database Administrator's Reference for UNIX* for more information about configuring swap space on AIX.
- 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 and platform
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirement 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 each system meets these requirements, follow these steps:

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

```
# /usr/sbin/lssattr -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/lspcs -a
```

If necessary, refer to 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 following command:

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

The expected output from this command 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, refer to 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 non-routable IP addresses for the private 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	143.47.43.104	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.

Check the Software Requirements

Refer to 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.

Checking the Software Requirements

Check that the required software and patches are installed on the system.

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:</p> <p> <code>bos.adt.base</code> <code>bos.adt.lib</code> <code>bos.adt.libm</code> <code>bos.perf.libperfstat</code> <code>bos.perf.perfstat</code> <code>bos.perf.proctools</code> <code>rsct.basic.rte</code> </p>
Real Application Clusters	<p>High Availability Cluster Multi-Processing (HACMP) v5.1</p> <p>Note: HACMP is required only if you want to use raw logical volumes for Oracle CRS or database file storage, however, it is supported for all installations.</p> <p>General Parallel File System (GPFS) v2.1 or v2.2</p> <p>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:</p> <p><code>x11.motif.lib</code></p> <p>Note: This software is required only to build the sample programs.</p>
Oracle Messaging Gateway	<p>IBM WebSphere MQ V5.3, client and server:</p> <p> <code>mqm.Client.Bnd</code> <code>mqm.Server.Bnd</code> </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>
PL/SQL native compilation, 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:

<http://www-912.ibm.com/eserver/support/fixes/>

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

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

If a fileset is not installed and committed, then install it. Refer to your operating system or software documentation for information about installing filesets.

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 READAHEAD/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 ■ IY56024: CIO WRITE RETURNS INCORRECT LENGTH <p>The following program technical fix (PTF) is required if you intend to use ASM for database file storage:</p> <p>U496549: bos.rte.aio.5.2.0.15</p>
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>PTFs required for VisualAge C compiler:</p> <p>U489726: vac.C.6.0.0.4 (or later)</p>
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 v2.1:</p> <ul style="list-style-type: none"> ■ IY52454: DirectIO fixes for Linux, Inerop: backup extended ■ U489058: mmfs.base.cmds.3.5.0.6 ■ U496347: mmfs.gpfs.rte.3.5.0.10 ■ U496395: mmfs.gpfs.rte.2.1.0.10 <p>APARs required for GPFS v2.2:</p> <ul style="list-style-type: none"> ■ IY54739: GPFS 2.2 mandatory service

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 "IY43980 IY44810 IY45462 IY45707 IY46214 IY46605 \
IY48525 IY51801 IY56024"
```

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

<http://www-912.ibm.com/eserver/support/fixes/>

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

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

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

<http://www-912.ibm.com/eserver/support/fixes/>

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

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

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 refer to 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 OSDBA 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. In the following procedure, use the group name `dba` unless a group with that name already exists.

1. Enter the following command:


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

Creating an OSOPER Group (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. In the following procedure, use 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. Refer to 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, refer to the "[Verifying that the UNIX User nobody Exists](#)" section on page 2-13.
- To modify an existing user, refer to the "[Modify an Existing Oracle Software Owner User](#)" section on page 2-13.
- To create a new user, refer to 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. Use 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
```

Refer to the "[Verifying that the UNIX User nobody Exists](#)" section on page 2-13 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 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 required 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
```

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 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`. Refer to 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 node1-public date
$ rsh node2-public 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 and Shell Limits

Refer to the following section for information about configuring the system kernel parameters and shell limits.

Configuring Shell Limits, System Configuration, and Network Tuning Parameters

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

Oracle recommends that you set shell limits, system configuration parameters, and network tuning parameters as described in this section 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:

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.

Configure Network Tuning Parameters

Verify that the network tuning parameters shown in the following table are set to the values shown or higher values. The procedure following the table describes how to verify and set the values.

Network Tuning Parameter	Recommended Value
ipqmaxlen	512
rfc1323	1
sb_max	2*655360
tcp_recvspace	65536
tcp_sendspace	65536
udp_recvspace	655360 Note: The recommended value of this parameter is 10 times the value of the udp_sendspace parameter. The value must be less than the value of the sb_max parameter.
udp_sendspace	65536 Note: This value is suitable for a default database installation. For production databases, the minimum value for this parameter is 4 KB plus the value of the database DB_BLOCK_SIZE initialization parameter multiplied by the value of the DB_MULTIBLOCK_READ_COUNT initialization parameter: $(DB_BLOCK_SIZE * DB_MULTIBLOCK_READ_COUNT) + 4 \text{ KB}$

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

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

```
# /usr/sbin/no -a | more
```

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

```
# /usr/sbin/lssattr -E -l sys0 -a pre520tune
```

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

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

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

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

```
# /usr/sbin/no -o parameter_name=value
```

For example:

```
# /usr/sbin/no -o udp_recvspace=655360
```

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

```
if [ -f /usr/sbin/no ] ; then
  /usr/sbin/no -o udp_sendspace=65536
  /usr/sbin/no -o udp_recvspace=655360
  /usr/sbin/no -o tcp_sendspace=65536
  /usr/sbin/no -o tcp_recvspace=65536
  /usr/sbin/no -o rfc1323=1
  /usr/sbin/no -o sb_max=2*655360
  /usr/sbin/no -o ipqmaxlen=512
fi
```

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

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

- `ipqmaxlen` parameter:

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

- Other parameter:

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

Note: If you modify the `ipqmaxlen` parameter, you must restart the system.

These commands modify the `/etc/tunables/nextboot` file, causing the attribute values to persist when the system restarts.

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 or an NFS file system on a NAS device, you must specify a path for the Oracle Inventory directory on a local file system. The Oracle base directory must be 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. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
/u01/crs/oracle/product/10.1.0/crs
```

Note: Because you must change the permissions of all of the parent directories of the CRS home directory after installing the software to grant write access only to the `root` user, the CRS home directory must not be a subdirectory of the Oracle base 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.

- To determine the free disk space on the file system where the Oracle base directory is located, enter the following command:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, refer to the "[Create the CRS Home Directory](#)" section on page 2-22.

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, refer to the following section.

Creating a New Oracle Base Directory

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/oracle_sw_owner
# chmod -R 775 /mount_point/app/oracle_sw_owner
```

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.

Create the CRS Home Directory

You must create a CRS home directory before installing Oracle CRS. You can create the CRS home directory on the same file system as the Oracle base directory, or you can choose a different file system for this directory. If you choose the same file system as the Oracle base directory, you must not use a subdirectory of the Oracle base directory as the CRS home directory.

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 at least 1 MB of free disk space.

If you are using the same file system for the Oracle base directory, this 1 MB of disk space is additional to the free disk space requirement that you identified previously.

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

The path to the CRS home 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 CRS home 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/crs/oracle_sw_owner/product/10.1.0/crs
# chown -R oracle:oinstall /mount_point/crs/oracle_sw_owner
# chmod -R 775 /mount_point/crs/oracle_sw_owner
```

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

```
/u01/crs/oracle/product/10.1.0/crs
```

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 installing Oracle CRS, set the `ORACLE_HOME` environment variable to specify this directory.

Note: During the installation, before you run the `root.sh` script, you must change the permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories.

Choose a Storage Option for Oracle CRS, Database, and Recovery Files

The following table shows the storage options supported for storing Oracle Cluster Ready Services (CRS) files, Oracle database files, and Oracle database recovery files. Oracle database files include datafiles, control files, redo log files, the server parameter file, and the password file. Oracle CRS files include the Oracle Cluster Registry (OCR) and the CRS voting disk.

For all installations, you must choose the storage option that you want to use for Oracle CRS files and Oracle database files. If you want to enable automated backups during the installation, you must also choose the storage option that you want to use for recovery files (the flash recovery area). You do not have to use the same storage option for each file type.

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

<http://metalink.oracle.com>

Storage Option	File Types Supported		
	CRS	Database	Recovery
Automatic Storage Management	No	Yes	Yes

Storage Option	File Types Supported		
	CRS	Database	Recovery
Cluster file system Note: Requires GPFS	Yes	Yes	Yes
Shared raw logical volumes Note: Requires HACMP	Yes	Yes	No
Shared raw disk devices	Yes	Yes	No

Use the following guidelines when choosing the storage options that you want to use for each file type:

- You can choose any combination of the supported storage options for each file type as long as you satisfy any requirements listed for the chosen storage options.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For Standard Edition installations, ASM is the only supported storage option for database or recovery files.
- You cannot use Automatic Storage Management to store Oracle CRS files, because these files must be accessible before any Oracle instance starts.
- If you are not using HACMP, you cannot use shared raw logical volumes for CRS or database file storage.

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

- To use a file system for Oracle CRS, database, or recovery file storage, refer to the ["Create Directories for Oracle CRS, Database, or Recovery Files"](#) section on page 2-24.
- To use ASM for database or recovery file storage, refer to the ["Configure Disks for Automatic Storage Management"](#) section on page 2-28.
- To use raw devices (disks or logical volumes) for Oracle CRS or database file storage, refer to the ["Configure Raw Devices"](#) section on page 2-34.

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 on shared disks, for example, a GPFS cluster file system.
- 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 that 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 on shared disks, for example, a GPFS cluster file system.
- 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 that 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 that 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 on shared disks, for example, a GPFS cluster file system.

- 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, refer to the *Oracle Backup and Recovery Basics* manual.

- The default path suggested by the Installer for the flash recovery area 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 that 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: The mount point that you use for the file system must be identical on each node. Make sure that the file systems are configured to mount automatically when a node restarts.

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
Database files	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, refer to one of the following sections:
 - [Configure Disks for Automatic Storage Management](#) on page 2-28
 - [Configure Raw Devices](#) on page 2-34

Otherwise, refer to the "[Verify the Cluster Software Configuration](#)" section on page 2-47.

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](#)
- [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. Refer to the *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, recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. You can use the file system for one file type and ASM for the other.

For RAC installations, if you choose to enable automated backups and you do not have a shared 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 on 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 database files 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 method 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 database files 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	Database Files	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

For RAC installations, 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 ASM instance is already running 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 protect against 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, refer to the "[Configuring Disks for ASM](#)" section on page 2-31.

Using an Existing ASM Disk Group

If you want to store either database or recovery files in an existing ASM disk group, 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 method 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 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 AIX, follow these steps:

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

- To make sure that the disks are available, enter the following command on every node:

```
# /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
```

- If a disk is not listed as available on any node, enter the following command to configure the new disks:

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

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

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

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

```
hdisk2      0000078752249812  None
```

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

- If a disk that you want to use for the disk group does not have a PVID, enter a command similar to the following to assign one to it:

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

- On each of the other nodes, enter a command similar to the following to identify the device name associated with each PVID on that node:

```
# /usr/sbin/lspv | grep -i 0000078752249812
```

Depending on how each node is configured, the device names might differ between nodes.

- To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute listed in the following table to the value shown, depending on the disk type:

Disk Type	Attribute	Value
SSA or FASTT disks	reserve_lock	no
ESS, EMC, HDS, CLARiiON, or MPIO-capable disks	reserve_policy	no_reserve

To determine whether the attribute has the correct value, enter a command similar to the following on all cluster nodes for each disk device that you want to use:

```
# /usr/sbin/lsattr -E -l hdiskn
```

If the required attribute is not set to the correct value on any node, enter a command similar to one of the following on that node:

- SSA and FASTT devices:

```
# /usr/sbin/chdev -l hdiskn -a reserve_lock=no
```

- ESS, EMC, HDS, CLARiiON, and MPIO-capable devices:

```
# /usr/sbin/chdev -l hdiskn -a reserve_policy=no_reserve
```

8. Enter commands similar to the following on any node to clear the PVID from each disk device that you want to use:

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

9. Enter commands similar to the following on every node 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, make sure that you set the permissions only on the correct logical device name for the disk.

The device name associated with a disk might be different on other nodes. Make sure that you specify the correct device name on each node.

10. If you also want to use raw devices for storage, refer to the "[Configure Raw Devices](#)" section on page 2-34.

Otherwise, refer to the "[Verify the Cluster Software Configuration](#)" section on page 2-47.

Configure Raw Devices

Refer to the following section for information about configuring raw devices.

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

Configuring Raw Disk Devices or Raw Logical Volumes

The following subsections describe how to configure raw disk devices or raw logical volumes on AIX:

- [Configuring Raw Disk Devices for Oracle CRS or Database File Storage for Installations Without HACMP or GPFS](#) on page 2-34
Follow the procedures described in this section if you are installing RAC on a cluster without HACMP or GPFS and you intend to use raw devices for CRS or database file storage.
- [Configuring Raw Logical Volumes for Oracle CRS or Database File Storage](#) on page 2-39
Follow the procedures described in this section only if you are using HACMP.

Configuring Raw Disk Devices for Oracle CRS or Database File Storage for Installations Without HACMP or GPFS

If you are installing RAC on an AIX cluster without HACMP or GPFS, you must use shared raw disk devices for the Cluster Ready Services (CRS) files. You can also use shared raw disk devices for database file storage, however, Oracle recommends that you use ASM to store database files in this situation. This section describes how to configure the shared raw disk devices for Oracle CRS files (Oracle Cluster Registry and CRS voting disk) and database files.

[Table 2–1](#) lists the number and size of the raw disk devices that you must configure for database files. [Table 2–2](#) lists the number and size of the raw disk devices that you must configure for CRS files.

Note: Because each file requires exclusive use of a complete disk device, Oracle recommends that, if possible, you use disk devices with sizes that closely match the size requirements of the files that they will store. You cannot use the disks that you choose for these files for any other purpose.

Table 2–1 Raw Disk Devices Required for Database Files on AIX

Number	Size (MB)	Purpose and Sample Alternative Device File Name
1	500	SYSTEM tablespace: <i>dbname_system_raw_500m</i>
1	300 + (Number of instances * 250)	SYSAUX tablespace: <i>dbname_sysaux_raw_800m</i>

Table 2–1 (Cont.) Raw Disk Devices Required for Database Files on AIX

Number	Size (MB)	Purpose and Sample Alternative Device File Name
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	250	TEMP tablespace: <i>dbname_temp_raw_250m</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	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwdfile_raw_5m</i>

Table 2–2 Raw Disk Devices Required for CRS Files on AIX

Number	Size (MB)	Purpose and Sample Alternative Device File Name
1	100	Oracle Cluster Registry: <i>ora_ocr_raw_100m</i> Note: You need to create this raw device only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle Cluster Registry (OCR).
1	20	Oracle CRS voting disk: <i>ora_vote_raw_20m</i> Note: You need to create this raw device 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 configure shared raw disk devices for CRS files, database files, or both, follow these steps:

1. If you intend to use raw disk devices for database file storage, 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. Identify or configure the required disk devices.

The disk devices must be shared on all of the cluster nodes.
3. As the `root` user, enter the following command on any node to identify the device names for the disk devices that you want to use:

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

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

```
hdisk17          0009005fb9c23648          None
```

In this example, `hdisk17` is the device name of the disk and `0009005fb9c23648` is the physical volume ID (PVID).

4. If a disk device that you want to use does not have a PVID, enter a command similar to the following to assign one to it:

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

5. On each of the other nodes, enter a command similar to the following to identify the device name associated with each PVID on that node:

```
# /usr/sbin/lspv | grep -i "0009005fb9c23648"
```

The output from this command should be similar to the following:

```
hdisk18          0009005fb9c23648          None
```

In this example, the device name associated with the disk device (`hdisk18`) is different on this node.

6. If the device names are the same on all nodes, enter commands similar to the following on all nodes to change the owner, group, and permissions on the character raw device files for the disk devices:

- OCR device:

```
# chown root:oinstall /dev/rhdiskn
# chmod 640 /dev/rhdiskn
```

- Other devices:

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

7. If the device name associated with the PVID for a disk that you want to use is different on any node, you must create a new device file for the disk on each of the nodes using a common unused name.

For the new device files, choose an alternative device file name that identifies the purpose of the disk device. The previous table suggests alternative device file names for each file. For database files, replace *dbname* in the alternative device file name with the name that you chose for the database in step 1.

Note: Alternatively, you could choose a name that contains a number that will never be used on any of the nodes, for example `hdisk99`.

To create a new common device file for a disk device on all nodes, follow these steps on each node:

- a. Enter the following command to determine the device major and minor numbers that identify the disk device, where *n* is the disk number for the disk device on this node:

```
# ls -alF /dev/*hdiskn
```


The output from this command is similar to the following:

```
brw----- 1 root system    24,8192 Dec 05 2001  /dev/hdiskn
crw----- 1 root system    24,8192 Dec 05 2001  /dev/rhdiskn
```

In this example, the device file `/dev/rhdiskn` represents the character raw device, 24 is the device major number, and 8192 is the device minor number.

- b. Enter a command similar to the following to create the new device file, specifying the new device file name and the device major and minor numbers that you identified in the previous step:

Note: In the following example, you must specify the character `c` to create a character raw device file.

```
# mknod /dev/ora_ocr_raw_100m c 24 8192
```

- c. Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for the disk:

– OCR:

```
# chown root:oinstall /dev/ora_ocr_raw_100m
# chmod 640 /dev/ora_ocr_raw_100m
```

– CRS voting disk or database files:

```
# chown oracle:dba /dev/ora_vote_raw_20m
# chmod 660 /dev/ora_vote_raw_20m
```

- d. Enter a command similar to the following to verify that you have created the new device file successfully:

```
# ls -alF /dev | grep "24,8192"
```

The output should be similar to the following:

```
brw----- 1 root  system    24,8192 Dec 05 2001  /dev/hdiskn
crw-r----- 1 root  oinstall  24,8192 Dec 05 2001  /dev/ora_ocr_raw_100m
crw----- 1 root  system    24,8192 Dec 05 2001  /dev/rhdiskn
```

8. To enable simultaneous access to a disk device from multiple nodes, you must set the appropriate Object Data Manager (ODM) attribute listed in the following table to the value shown, depending on the disk type:

Disk Type	Attribute	Value
SSA or FASTt disks	reserve_lock	no
ESS, EMC, HDS, CLARiiON, or MPIO-capable disks	reserve_policy	no_reserve

To determine whether the attribute has the correct value, enter a command similar to the following on all cluster nodes for each disk device that you want to use:

```
# /usr/sbin/lsattr -E -l hdiskn
```

If the required attribute is not set to the correct value on any node, enter a command similar to one of the following on that node:

- SSA and FASTT devices


```
# /usr/sbin/chdev -l hdiskn -a reserve_lock=no
```
 - ESS, EMC, HDS, CLARiiON, and MPIO-capable devices


```
# /usr/sbin/chdev -l hdiskn -a reserve_policy=no_reserve
```
9. Enter commands similar to the following on any node to clear the PVID from each disk device that you want to use:
- ```
/usr/sbin/chdev -l hdiskn -a pv=clear
```
10. If you are using raw disk devices for database files, follow these steps to create the DBCA raw device mapping file:

---

**Note:** You must complete this procedure only if you are using raw devices for database files. The DBCA raw device mapping file enables the DBCA to identify the appropriate raw disk device for each database file. You do not specify the raw devices for the Oracle CRS files in the DBCA raw device mapping file.

---

- a. 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
```
- b. 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.

- c. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
- d. Using any text editor, create a text file similar to the following that identifies the disk device file name associated with each database file.

Oracle recommends that you use a file name similar to `dbname_raw.conf` for this file.

Note: The following example shows a sample mapping file for a two-instance RAC cluster. Some of the devices use alternative disk device file names. Make sure that the device file name that you specify identifies the same disk device on all nodes.

```
system=/dev/rhdisk11
sysaux=/dev/rhdisk12
example=/dev/rhdisk13
```

```

users=/dev/rhdisk14
temp=/dev/rhdisk15
undotbs1=/dev/rhdisk16
undotbs2=/dev/rhdisk17
redo1_1=/dev/rhdisk18
redo1_2=/dev/rhdisk19
redo2_1=/dev/rhdisk20
redo2_2=/dev/rhdisk22
control1=/dev/rhdisk23
control2=/dev/rhdisk24
spfile=/dev/dbname_spfile_raw_5m
pwdfile=/dev/dbname_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=device_file_name
```

The alternative device file names suggested in the previous table include the database object identifier that you must use in this mapping file. For example, in the following alternative disk device name, *redo1_1* is the database object identifier:

```
/dev/rac_redo1_1_raw_120m
```

- 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 tablespace datafiles.
- e. Save the file and note the file name that you specified.
 - f. 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.
11. When you are installing Oracle Cluster Ready Services, you must enter the paths to the appropriate device files when prompted for the path of the OCR and CRS voting disk, for example:

```
/dev/rhdisk10
```

12. To continue the installation, refer to the ["Verify the Cluster Software Configuration"](#) section on page 2-47

Configuring Raw Logical Volumes for Oracle CRS or Database File Storage

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

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, refer to the HACMP documentation.

Note: If you are upgrading a database, you must also create a new logical volume for the SYSAUX tablespace. Refer to the "[Create Raw Logical Volumes in the New Volume Group](#)" section on page 2-42 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, refer to 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 on every node:

```
# /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 on any node, enter the following command to configure the new disks:

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

4. Enter the following command on any node to identify the device names and any associated volume group for each disk:

```
# /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 a disk that you want to use for the volume group does not have a PVID, enter a command similar to the following to assign one to it:

```
# /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.

- 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. Refer to 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 that 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).
-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.

- 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. Identify the logical volumes that you must create.

Table 2–3 lists the number and size of the logical volumes that you must create for database files. Table 2–4 lists the number and size of the logical volumes that you must create for CRS files.

Table 2–3 Raw Logical Volumes Required for Database Files on AIX

Number	Size (MB)	Purpose and Sample Logical Volume Name
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	250	TEMP tablespace: <code>dbname_temp_raw_250m</code>
1	160	EXAMPLE tablespace: <code>dbname_example_raw_160m</code>
1	120	USERS tablespace: <code>dbname_users_raw_120m</code>
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): <code>dbname_redonm_raw_120m</code>
2	110	First and second control files: <code>dbname_control{1 2}_raw_110m</code>
1	5	Server parameter file (SPFILE): <code>dbname_spfile_raw_5m</code>
1	5	Password file: <code>dbname_pwdfile_raw_5m</code>

Table 2–4 Raw Logical Volumes Required for CRS Files on AIX

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	100	Oracle Cluster Registry: <code>ora_ocr_raw_100m</code> 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 (OCR). 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.

Table 2–4 (Cont.) Raw Logical Volumes Required for CRS Files on AIX

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	20	Oracle CRS voting disk: ora_vote_raw_20m 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_raw_800m -T 0 -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:oinstall /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 a command 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:oinstall /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/rdbname_system_raw_500m
sysaux=/dev/rdbname_sysaux_raw_800m
example=/dev/rdbname_example_raw_160m
users=/dev/rdbname_users_raw_120m
temp=/dev/rdbname_temp_raw_250m
undotbs1=/dev/rdbname_undotbs1_raw_500m
undotbs2=/dev/rdbname_undotbs2_raw_500m
redo1_1=/dev/rdbname_redo1_1_raw_120m
```

```
redo1_2=/dev/rdbname_redo1_2_raw_120m
redo2_1=/dev/rdbname_redo2_1_raw_120m
redo2_2=/dev/rdbname_redo2_2_raw_120m
control1=/dev/rdbname_control1_raw_110m
control2=/dev/rdbname_control2_raw_110m
spfile=/dev/rdbname_spfile_raw_5m
pwdfile=/dev/rdbname_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/rac_redo1_1_raw_120m
```

- 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 tablespace datafiles.
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

The following sections describe how to verify that the cluster software is configured correctly:

- [Verify that the HACMP Cluster Software is Configured Correctly](#)
- [Synchronize the System Time on Cluster Nodes](#)

Verify that the HACMP Cluster Software is Configured Correctly

Note: HACMP is not required. However, if it is installed, Oracle 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.

Synchronize the System Time on Cluster Nodes

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:

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

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

2. Using any text editor, edit the `/etc/ntp.conf` file:

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

3. 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
```

4. To start `xntpd`, follow these steps:

- a. Enter the following command:

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

- b. 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.

Note: If no Oracle Net listeners are running, refer to the ["Configure the oracle User's Environment"](#) section on page 2-50 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 (`csh` or `tcsh`):

```
% vi .login
```

7. Enter or edit the following line, specifying a value of 022 for the default file mode 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. Add one of the following lines to the file to set the `AIXTHREAD_SCOPE` environment variable to `S` (system-wide thread scope):

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

```
AIXTHREAD_SCOPE=S; export AIXTHREAD_SCOPE
```

- C shell (csh or tcsh):

```
setenv AIXTHREAD_SCOPE S
```

Note: You must set this environment variable in the `oracle` user's shell startup file on all cluster nodes. For more information about this environment variable, refer to the *Oracle Database Administrator's Reference for UNIX Systems*.

10. Save the file and exit from the editor.

11. To run the shell startup script, enter one of the following commands:

- Bourne, Bash, or Korn shell:

```
$ . ./profile
```

- C shell:

```
% source ./login
```

12. 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 that you want to use to display the Installer (your workstation or PC).

13. 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:

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

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

```
$ su - root
# mkdir /mount_point/tmp
# chmod a+wr /mount_point/tmp
# exit
```

- 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
```

14. 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 or identified earlier.

15. If you are using raw devices for database file 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
```

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

■ Bourne, Bash, or Korn shell:

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

■ C shell:

```
% unsetenv ORACLE_HOME
% unsetenv TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value that 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.

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

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


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

Pre-Installation Tasks for RAC on HP-UX

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer (OUI). 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](#)
- [Create the CRS Home Directory](#)
- [Choose a Storage Option for Oracle CRS, Database, and Recovery Files](#)
- [Configure Disks for Automatic Storage Management](#)
- [Configure Raw Devices](#)
- [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-mode 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 silent-mode installations, refer to [Appendix B](#).

- 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 are not installing the software on the local system, use the `ssh`, `rlogin`, or `telnet` commands to connect to the system where you want to install the software:

```
$ 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, refer to 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

Each 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 and platform
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirement 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 each system meets these requirements, follow these steps:

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

```
# grep "Physical:" /var/adm/syslog/syslog.log
```

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, refer to 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 following command:

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

The expected output from this command is the value 64.

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, refer to 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)
 - (PA-RISC only) Hyper Messaging protocol (HMP) using HyperFabric adapters and switches

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

To use HMP as the interconnect protocol for RAC, you must relink the `oracle` executable. Refer to the *Oracle Database Administrator's Reference for UNIX Systems* for more information about enabling HMP for RAC on HP-UX 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 non-routable IP addresses for the private 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	143.47.43.104	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

Refer to 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.

Checking the Software Requirements

Check that the required software and patches are installed on the system.

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: <ul style="list-style-type: none"> ■ PA-RISC (64-bit) systems: HP-UX 11i (11.11) PA-RISC ■ Itanium systems: HP-UX 11i v2 (11.23)
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	Note: HP Serviceguard is required only if you want to use shared logical volumes for CRS or database files. HP-UX 11i (11.11): <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 HP-UX 11i v2 (11.23): <ul style="list-style-type: none"> ■ HP Serviceguard A11.15.01 or higher ■ HP Serviceguard Extension for RAC A11.15.01 or higher
Oracle Messaging Gateway (PA-RISC only)	Oracle Messaging Gateway supports the integration of Oracle Streams Advanced Queuing (AQ) with the following software: <ul style="list-style-type: none"> ■ IBM WebSphere MQ V5.3, client and server: MQSERIES.MQM-CL-HPUX MQSERIES.MQM-SERVER ■ TIBCO Rendezvous 7.2
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	HP-UX 11i (11.11): HP C/ANSI C Compiler (B.11.11.08): C-ANSI-C HP aC++ Compiler (C.03.50): ACXX HP-UX 11i v2 (11.23): HP C/ANSI C Compiler (C.05.50): C-ANSI-C HP aC++ Compiler (C.05.50): ACXX

Installation Type or Product	Requirement
Oracle JDBC/OCI Drivers	<p>You can use the following optional Java SDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation:</p> <p>HP-UX 11i (11.11):</p> <ul style="list-style-type: none"> ■ Java SDK 1.3.1.02 with the JNDI extension ■ Java SDK 1.2.2.09 <p>HP-UX 11i v2 (11.23):</p> <ul style="list-style-type: none"> ■ Java SDK 1.3.1.09 <p>Note: Java SDK 1.4.2.00 (PA-RISC) or 1.4.1.03 (Itanium) is installed with this release.</p>
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE (PA-RISC only)	<p>DCE (B.11.11.10):</p> <ul style="list-style-type: none"> DCE-C-Tools DCE-CDS-Server DCE-Core DCE-CoreAdmin DCE-CoreTools DCE-SEC-Server
Oracle Transparent Gateway for Sybase (PA-RISC only)	<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>
Oracle Transparent Gateway for Informix (PA-RISC only)	<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 Ingres (PA-RISC only)	Ingres II version 2.0, 2.5, or 2.6
Oracle Transparent Gateway for Teradata (PA-RISC only)	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 fileset 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 fileset is not installed, you must install it. Refer to 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 (11.11): HP-UX 11i Quality Pack (GOLDQPK11i), June 2003 or later: GOLDAPPS11i GOLDBASE11i HP-UX 11i v2 (11.23): None currently required
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) Patches for HP-UX 11i v2 (11.23): <ul style="list-style-type: none"> ■ PHSS_29658: Aries cumulative patch ■ PHSS_29660: linker + fdp cumulative patch
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: Refer to the following Web site for information about additional patches that might be required by JDK 1.4.2:</p> <p>http://www.hp.com/products1/unix/java/patches/index.html</p>

Installation Type or Product	Requirement
Real Application Clusters	<p>Patches for HP-UX 11i (11.11): PHSS_29096: ServiceGuard Extension for RAC 11.15.00</p> <p>Patches for HP-UX 11i v2 (11.23):</p> <ul style="list-style-type: none"> ■ PHSS_29097: ServiceGuard Extension for RAC A.11.15 ■ PHSS_30088: MC/ServiceGuard A.11.15.00
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Patches for HP-UX 11i (11.11):</p> <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 <p>Patches for HP-UX 11i v2 (11.23):</p> <ul style="list-style-type: none"> ■ PHSS_29655: aC++ Compiler (A.05.52) ■ PHSS_29656: HP C Compiler (A.05.52) ■ PHSS_29657: u2comp/be/plugin library Patch
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE (PA-RISC only)	<p>Patches for HP-UX 11i (11.11):</p> <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 (PA-RISC only)	<p>Corrective service diskettes (CSDs) for WebSphere MQ:</p> <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. On PA-RISC systems only, enter the following command to determine whether the HP-UX 11i Quality Pack is installed:

```
# /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, refer to 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

On HP-UX systems, 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 refer to 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, enter the following command to create the `oinstall` group:

```
# /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 OSDBA 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, enter the following command to create the `dba` group. Use 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, enter the following command to create the `oper` group. Use 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. Refer to 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, refer to the ["Creating an Unprivileged User"](#) section on page 3-14.
- To modify an existing user, refer to the [Modify an Existing Oracle Software Owner User](#) section on page 3-14.
- To create a new user, refer to 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. Use 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
```

Refer to the ["Creating an Unprivileged User"](#) section on page 3-14 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, enter a command similar to the following to modify it. Specify 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
```

Creating an Unprivileged User

You must create a new unprivileged local user to own the `extjob` executable. You can use any user name for this user, but this document uses the name `extjob`. To create the `extjob` 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 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`. Refer to 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 node1-public date
$ remsh node2-public 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 on HP-UX systems, you must grant several system privileges to the OSDBA group. Complete the following steps if you are installing the software on HP-UX:

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 restarts:

```
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/bin/getprivgrp dba
```

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

Configure Kernel Parameters and Shell Limits

Refer to the following section for information about configuring the system kernel parameters and shell limits.

Configuring 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	1024
maxdsiz	1073741824 (1 GB)
maxdsiz_64bit	2147483648 (2 GB)
maxssiz	134217728 (128 MB)
maxssiz_64bit	1073741824 (1 GB)
maxswapchunks (PA-RISC only)	16384
maxuprc	$((nproc*9)/10)$
msgmap	$(msgtql+2)$
msgmni	$(nproc)$
msgseg	32767
msgtql	$(nproc)$
ncsize (PA-RISC)	$(ninode+vx_ncsize)$
ncsize (Itanium)	$(ninode+1024)$
nfile	$(15*nproc+2048)$
nflocks	$(nproc)$
ninode	$(8*nproc+2048)$
nkthread	$((nproc*7)/4)+16)$
nproc	4096
semmap (PA-RISC only)	$(semmni+2)$
semmni	$(nproc)$
semmns	$(semmni*2)$
semmnu	$(nproc-4)$
semvmx	32767
shmmax	The size of physical memory or 1073741824 (0X40000000), 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 on PA-RISC systems, 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:

1. Follow these steps, depending on your platform:

PA-RISC:

- a. 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
```

- b. Start System Administration Manager (SAM):

```
# /usr/sbin/sam
```

- c. Choose the **Kernel Configuration** area, then choose the **Configurable Parameters** area.

- d. Check the value or formula specified for each of these parameters and, if necessary, modify that value or formula.

If necessary, refer to the SAM online help for more information about completing this step.

Note: If you modify the value of a parameter that is not dynamic, you must restart the system.

Itanium:

- a. Enter the following command to start the kcweb application:

```
# /usr/sbin/kcweb -F
```

- b. Check the value or formula specified for each of these parameters and, if necessary, modify that value or formula.

If necessary, refer to the kcweb online help for more information about completing this step.

Note: If you modify the value of a parameter that is not dynamic, you must restart the system.

2. If necessary, when the system restarts, log in and switch user to root.
3. 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
```

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. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
/u01/crs/oracle/product/10.1.0/crs
```

Note: Because you must change the permissions of all of the parent directories of the CRS home directory after installing the software to grant write access only to the `root` user, the CRS home directory must not be a subdirectory of the Oracle base 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 3 GB of free disk space.

To determine the free disk space on the file system where the Oracle base directory is located, enter the following command:

```
# bdf oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, refer to the ["Create the CRS Home Directory"](#) 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, refer to 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 3 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/oracle_sw_owner
# chmod -R 775 /mount_point/app/oracle_sw_owner
```

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.

Create the CRS Home Directory

You must create a CRS home directory before installing Oracle CRS. You can create the CRS home directory on the same file system as the Oracle base directory, or you can choose a different file system for this directory. If you choose the same file system as

the Oracle base directory, you must not use a subdirectory of the Oracle base directory as the CRS home directory.

To identify an appropriate file system, follow these steps:

1. Use the `df` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has at least 1 MB of free disk space.

If you are using the same file system for the Oracle base directory, this 1 MB of disk space is additional to the free disk space requirement that you identified previously.

Note: The file system must be a local file system.

The path to the CRS home 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 CRS home 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/crs/oracle_sw_owner/product/10.1.0/crs
# chown -R oracle:oinstall /mount_point/crs/oracle_sw_owner
# chmod -R 775 /mount_point/crs/oracle_sw_owner
```

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

```
/u01/crs/oracle/product/10.1.0/crs
```

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 installing Oracle CRS, set the `ORACLE_HOME` environment variable to specify this directory.

Note: During the installation, before you run the `root.sh` script, you must change the permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories.

Choose a Storage Option for Oracle CRS, Database, and Recovery Files

The following table shows the storage options supported for storing Oracle Cluster Ready Services (CRS) files, Oracle database files, and Oracle database recovery files. Oracle database files include datafiles, control files, redo log files, the server parameter file, and the password file. Oracle CRS files include the Oracle Cluster Registry (OCR) and the CRS voting disk.

For all installations, you must choose the storage option that you want to use for Oracle CRS files and Oracle database files. If you want to enable automated backups during the installation, you must also choose the storage option that you want to use

for recovery files (the flash recovery area). You do not have to use the same storage option for each file type.

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

<http://metalink.oracle.com>

Storage Option	File Types Supported		
	CRS	Database	Recovery
Automatic Storage Management	No	Yes	Yes
Shared raw logical volumes	Yes	Yes	No
Note: Requires HP Serviceguard Extension for RAC and LVM			
Shared raw disk devices	Yes	Yes	No

Use the following guidelines when choosing the storage options that you want to use for each file type:

- You can choose any combination of the supported storage options for each file type as long as you satisfy any requirements listed for the chosen storage options.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For Standard Edition installations, ASM is the only supported storage option for database or recovery files.
- You cannot use Automatic Storage Management to store Oracle CRS files, because these files must be accessible before any Oracle instance starts.
- If you are not using HP Serviceguard Extension for RAC, you cannot use shared raw logical volumes for CRS or database file storage.

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

- To use ASM for database or recovery file storage, refer to the "[Configure Disks for Automatic Storage Management](#)" section on page 3-24.
- To use raw devices (disks or logical volumes) for Oracle CRS or database file storage, refer to the "[Configure Raw Devices](#)" section on page 3-29.

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](#)
- [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. Refer to the *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, recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. You can use raw devices for database files and ASM for recovery files if you choose.

For RAC installations, if you choose to enable automated backups, 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 on 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 database files 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 method 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 database files 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	Database Files	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

For RAC installations, 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 ASM instance is already running 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 protect against 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, refer to the "[Configuring Disks for ASM](#)" section on page 3-28.

Using an Existing ASM Disk Group

If you want to store either database or recovery files in an existing ASM disk group, 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 method 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 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 HP-UX, 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 on every node:

```
# /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/rdsk/cxydz`).

3. If the `iostats` command does not display device name information for a device that you want to use, enter the following command to install the special device files for any new devices:

```
# /usr/sbin/insf -e
```

4. For each disk that you want to add to a disk group, enter the following command on any node 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, refer to the appropriate documentation for information about verifying that a disk is not in use.

5. Enter commands similar to the following on every node 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/rdsk/cxydz
# chmod 660 /dev/rdsk/cxydz
```

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

If the nodes are configured differently, the device name for a particular device might be different on some nodes. Make sure that you specify the correct device names on each node.

6. If you also want to use raw devices for storage, refer to the "[Configure Raw Devices](#)" section on page 3-29.

Otherwise, refer to the "[Verify the Cluster Software Configuration](#)" section on page 3-42.

Configure Raw Devices

Refer to the following section for information about configuring raw devices.

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

Configuring Raw Disk Devices or Raw Logical Volumes

The following subsections describe how to configure raw disk devices or raw logical volumes on HP-UX:

- [Configuring Raw Disk Devices for Oracle CRS or Database File Storage for Installations Without HP Serviceguard Extension for RAC](#) on page 3-30

Follow the procedures described in this section if you are installing RAC on a cluster without HP Serviceguard Extension for RAC and you intend to use raw disk devices for CRS or database file storage.
- ["Configuring Shared Raw Logical Volumes for Oracle CRS or Database File Storage"](#) on page 3-35

Follow the procedures described in this section only if you are using HP Serviceguard Extension for RAC.

Configuring Raw Disk Devices for Oracle CRS or Database File Storage for Installations Without HP Serviceguard Extension for RAC

If you are installing RAC on an HP-UX cluster without HP Serviceguard Extension for RAC, you must use shared raw disk devices for the Cluster Ready Services (CRS) files. You can also use shared raw disk devices for database file storage, however, Oracle recommends that you use ASM to store database files in this situation. This section describes how to configure the shared raw disk devices for CRS files (Oracle Cluster Registry and CRS voting disk) and database files.

[Table 3–1](#) lists the number and size of the raw disk devices that you must configure for database files. [Table 3–2](#) lists the number and size of the raw disk devices that you must configure for CRS files.

Note: Because each file requires exclusive use of a complete disk device, Oracle recommends that, if possible, you use disk devices with sizes that closely match the size requirements of the files that they will store. You cannot use the disks that you choose for these files for any other purpose.

Table 3–1 Raw Disk Devices Required for Database Files on HP-UX

Number	Size (MB)	Purpose and Sample Alternative Device File Name
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>

Table 3–1 (Cont.) Raw Disk Devices Required for Database Files on HP-UX

Number	Size (MB)	Purpose and Sample Alternative Device File Name
1	250	TEMP tablespace: <i>dbname_temp_raw_250m</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 <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	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwdfile_raw_5m</i>

Table 3–2 Raw Disk Devices Required for CRS Files on HP-UX

Number	Size (MB)	Purpose and Sample Alternative Device File Name
1	100	Oracle Cluster Registry: <i>ora_ocr_raw_100m</i> Note: You need to create this raw device only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle Cluster Registry (OCR).
1	20	Oracle CRS voting disk: <i>ora_vote_raw_20m</i> Note: You need to create this raw device 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 configure shared raw disk devices for CRS files, database files, or both, follow these steps:

1. If you intend to use raw disk devices for database file storage, 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. Identify or configure the required disk devices.

The disk devices must be shared on all of the cluster nodes.

3. To make sure that the disks are available, enter the following command on every node:

```
# /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).

4. If the `ioscan` command does not display device name information for a device that you want to use, enter the following command to install the special device files for any new devices:

```
# /usr/sbin/insf -e
```

5. For each disk that you want to use, enter the following command on any node 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, refer to the appropriate documentation for information about verifying that a disk is not in use.

6. If the `ioscan` command shows different device names for the same device on any node, follow these steps on every node:

- a. Change directory to the /dev/rdisk directory.
- b. Enter the following command to list the raw disk device names and their associated major and minor numbers:

```
# ls -la
```

The output from this command is similar to the following for each disk device:

```
crw-r--r--  1 bin      sys      188 0x032000 Nov  4  2003 c3t2d0
```

In this example, 188 is the device major number and 0x32000 is the device minor number.

- c. Enter the following command to create a new device file for the disk that you want to use, specifying the same major and minor number as the existing device file:

Note: Oracle recommends that you use the alternative device file names shown in the previous table.

```
# mknod ora_ocr_raw_100m c 188 0x032000
```

- d. Repeat these steps on each node, specifying the correct major and minor numbers for the new device files on each node.
7. Enter commands similar to the following on every node to change the owner, group, and permissions on the character raw device file for each disk device that you want to use:

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

If you created an alternative device file for the device, set the permissions on that device file.

- OCR:

```
# chown root:oinstall /dev/rdisk/cxydz
# chmod 640 /dev/rdisk/cxydz
```

- CRS voting disk or database files:

```
# chown oracle:dba /dev/rdisk/cxydz
# chmod 660 /dev/rdisk/cxydz
```

8. If you are using raw disk devices for database files, follow these steps to create the DBCA raw device mapping file:

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

- a. 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
```

- b. 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.

- c. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
- d. Using any text editor, create a text file similar to the following that identifies the disk device file name associated with each database file.

Oracle recommends that you use a file name similar to `dbname_raw.conf` for this file.

Note: The following example shows a sample mapping file for a two-instance RAC cluster. Some of the devices use alternative disk device file names. Make sure that the device file name that you specify identifies the same disk device on all nodes.

```
system=/dev/rdisk/c2t1d1
sysaux=/dev/rdisk/c2t1d2
example=/dev/rdisk/c2t1d3
users=/dev/rdisk/c2t1d4
temp=/dev/rdisk/c2t1d5
undotbs1=/dev/rdisk/c2t1d6
undotbs2=/dev/rdisk/c2t1d7
redo1_1=/dev/rdisk/c2t1d8
redo1_2=/dev/rdisk/c2t1d9
redo2_1=/dev/rdisk/c2t1d10
redo2_2=/dev/rdisk/c2t1d11
control1=/dev/rdisk/c2t1d12
control2=/dev/rdisk/c2t1d13
spfile=/dev/rdisk/dbname_spfile_raw_5m
pwdfile=/dev/rdisk/dbname_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=device_file_name
```

The alternative device file names suggested in the previous table include the database object identifier that you must use in this mapping file. For example, in the following alternative disk device file name, `redo1_1` is the database object identifier:

```
rac_redo1_1_raw_120m
```

- 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 tablespace datafiles.
- e. Save the file and note the file name that you specified.
 - f. 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.
9. When you are installing Oracle Cluster Ready Services, you must enter the paths to the appropriate device files when prompted for the path of the OCR and CRS voting disk, for example:

```
/dev/rdisk/cxydz
```

10. To continue the installation, refer to the "[Verify the Cluster Software Configuration](#)" section on page 3-42.

Configuring Shared Raw Logical Volumes for Oracle CRS or Database File Storage

Note: 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. Refer to the HP-UX documentation for more information about using SAM.

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 Extension for RAC must be installed on all cluster nodes. If HP Serviceguard Extension for RAC is not installed, you can use shared raw disk devices to store the Oracle CRS or database 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 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, 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, refer to the HP Serviceguard or HP Serviceguard Extension for RAC documentation.

Note: If you are upgrading a database, you must also create a new logical volume for the SYSAUX tablespace. Refer to the "[Create Raw Logical Volumes in the New Volume Group](#)" section on page 3-38 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, refer to 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. Refer to 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 on every node, 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 restarts.

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 on every node:

```
# /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`).

3. If the `ioscan` command does not display device name information for a device that you want to use, enter the following command to install the special device files for any new devices:

```
# /usr/sbin/insf -e
```

4. For each disk that you want to add to the volume group, enter the following command on any node 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.

5. For each disk that you want to add to the volume group, enter a command similar to the following on any node:

```
# /sbin/pvcreate /dev/rdisk/cxydz
```

- 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 that you want to use for the volume group:

```
# mkdir /dev/vg_name
```

- 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 0x*nn*0000, 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.

- Identify an appropriate minor number that is unused on all nodes in the cluster.
- 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 that you want to give to the volume group
- 0x*nn*0000 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:

- 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*.

- Identify the logical volumes that you must create.

Table 3–3 lists the number and size of the logical volumes that you must create for database files. Table 3–4 lists the number and size of the logical volumes that you must create for CRS files.

Table 3–3 Raw Logical Volumes Required for Database Files on HP-UX

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	500	SYSTEM tablespace: <i>dbname_system_raw_500m</i>

Table 3–3 (Cont.) Raw Logical Volumes Required for Database Files on HP-UX

Number	Size (MB)	Purpose and Sample Logical Volume Name
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	250	TEMP tablespace: <i>dbname_temp_raw_250m</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	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwdfile_raw_5m</i>

Table 3–4 Raw Logical Volumes Required for CRS Files on HP-UX

Number	Size (MB)	Purpose and Sample Logical Volume Name
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 (OCR). 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:

```
# /sbin/lvcreate -n LVname -L size /dev/vg_name
```

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 that you chose for the database in step 1.

- *vg_name* 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_raw_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:oinstall /dev/vg_name/rora_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:oinstall /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 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 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 tablespace datafiles.
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: HP Serviceguard is not required. However, if it is installed, Oracle 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:

```
# /usr/sbin/cmruncl
```

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, refer to 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.

Note: If no Oracle Net listeners are running, refer to the ["Configure the oracle User's Environment"](#) section on page 3-44 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 (`csh` or `tcsh`):


```
% vi .login
```
7. Enter or edit the following line, specifying a value of 022 for the default file mode 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 that 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` command to identify a suitable file system with sufficient free space.
 - b. If necessary, enter commands similar to the following to create a temporary directory on the file system that you identified, and set the appropriate permissions on the directory:


```
$ 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 or identified earlier.

14. If you are using raw devices for database file 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:

■ Bourne, Bash, or Korn shell:

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

■ C shell:

```
% unsetenv ORACLE_HOME
% unsetenv TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value that 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 that the environment variables you set in this section have the correct values.

Pre-Installation Tasks for RAC on hp Tru64 UNIX

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer (OUI). 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 Subsystem Attributes](#)
- [Identify Required Software Directories](#)
- [Identify or Create an Oracle Base Directory](#)
- [Create the CRS Home Directory](#)
- [Choose a Storage Option for Oracle CRS, Database, and Recovery 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-mode 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 silent-mode installations, refer to [Appendix B](#).

- 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 are not installing the software on the local system, use the `ssh`, `rlogin`, or `telnet` commands to connect to the system where you want to install the software:


```
$ 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, refer to 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

Each 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 and platform

- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirement 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 each 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, refer to 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, refer to 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. Refer to 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 non-routable IP addresses for the private 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	143.47.43.104	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:

```
# /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

Refer to 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.

Checking the Software Requirements

Check that the required software and patches are installed on the system.

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:</p> <p>HP Tru64 UNIX V5.1B</p> <hr/> <p>Software Development Kit (SDK) v 1.4.2 for the Tru64 UNIX Operating System for the Java Platform (JDK 1.4.2)</p> <hr/> <p>Operating system subsets:</p> <p>OSFCMPLRS OSFLIBA OSFPGMR OSFSER OSFX11</p>
Oracle Real Application Clusters	<p>HP TruCluster 5.1B:</p> <p>TCRBASE TCRMIGRATE</p>
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE	<p>DCE v4.3.1 or later:</p> <p>DCECDS DECRTS DCESEC</p>
Oracle Spatial	<p>X Window and X/Motif software:</p> <p>OSFXDEVX OSFXINCLUDE OSFXLIBA</p> <p>Note: This software is required only to build the sample programs.</p>
Oracle Messaging Gateway	<p>IBM MQSeries V5.1 client and server:</p> <p>MQS_CLIENT MQS_SERVER</p> <p>MQSeries classes for Java and MQSeries classes for Java Message Service V5.2.2 (SupportPac MA88):</p> <p>MQS_MQJAVA</p>
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Compaq C Compiler V6.5-207 (dtk):</p> <p>DTCCMPLR</p> <p>Compaq C++ Version 6.5-014:</p> <p>CXXBASE CXXLIB CXXOLD</p>
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"> ■ SDK v 1.3.1-5 ■ SDK v 1.2.2-12 <p>Note: SDK v 1.4.2 must be installed before you install the software.</p>

Installation Type or Product	Requirement
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 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, refer to 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 Java SDK 1.4.2-3 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.htm
1
```

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>Tru64 UNIX V5.1B Patch Kit 2 or higher: T64V51BB22AS0002-20030415</p> <p>The following patch kits are also required:</p> <ul style="list-style-type: none"> ■ HP Tru64 UNIX 5.1B PK2 BL22 Fixes for AdvFS Panic in _OtsMove; and Possible Memory Corruption: T64KIT0020879-V51BB22-E-20031125 ■ HP Tru64 UNIX - Problem with IP Multicast Packets: T64KIT0019662-V51BB22-E-20030818 ■ HP Tru64 UNIX V5.1B PK2 (BL22) ERP Kit - Fix for Multiple applications may be granted an exclusive lock on the same file: T64KIT0021665-V51BB22-E-20040220 ■ Tru64 UNIX V5.1B PK2/BL22 Early Release Patch - Fix for Potential Application Core Dump: T64KIT0021681-V51BB22-E-20040223
Oracle Messaging Gateway	<p>Corrective service diskette (CSD) for MQSeries: CSD09 or later for MQSeries V5.1</p>

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, refer to 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 refer to 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, enter the following command to create the `oinstall` group:

```
# /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 OSDBA 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, enter the following command to create the `dba` group. Use 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, enter the following command to create the `oper` group. Use 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. Refer to 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, refer to the [Verifying that the UNIX User nobody Exists](#) section on page 4-12.
- To modify an existing user, refer to the "[Modify an Existing Oracle Software Owner User](#)" section on page 4-12.
- To create a new user, refer to 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. Use 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
```

Refer to the "[Verifying that the UNIX User nobody Exists](#)" section on page 4-12 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, enter a command similar to the following to modify it. Specify 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 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`. Refer to 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 node1-public date
$ rsh node2-public 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

Refer to the following section for information about configuring the kernel subsystem attributes.

Configuring Kernel Subsystem Attributes

Note: The kernel subsystem attribute 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. Refer to your operating system documentation for more information about tuning kernel subsystem attributes.

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 <code>rad_gh_regions[n]</code> or <code>gh_chunks</code> attributes are set in the <code>vm</code> subsystem. Otherwise, do not change the value.
proc	exec_disable_arg_limit	1
	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.

Subsystem	Attribute	Recommended Value
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	8193
vfs	fifo_do_adaptive	0
vm	new_wire_method	0

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 with a recommended value of 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:
    exec_disable_arg_limit = 1
    per_proc_stack_size = 8388608
    max_per_proc_stack_size = 33554432
    per_proc_data_size = 335544320
    max_per_proc_data_size = 335544320
```

```
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 restart 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, 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. Oracle recommends that you specify a path similar to the following for the CRS home directory:

`/u01/crs/oracle/product/10.1.0/crs`

Note: Because you must change the permissions of all of the parent directories of the CRS home directory after installing the software to grant write access only to the `root` user, the CRS home directory must not be a subdirectory of the Oracle base 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 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 3 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 4 GB

To determine the free disk space on the file system where the Oracle base directory is located, enter the following command:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, refer to the ["Create the CRS Home Directory"](#) 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, refer to 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 3 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 4 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/oracle_sw_owner
# chmod -R 775 /mount_point/app/oracle_sw_owner
```

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.

Create the CRS Home Directory

You must create a CRS home directory before installing Oracle CRS. You can create the CRS home directory on the same file system as the Oracle base directory, or you can choose a different file system for this directory. If you choose the same file system as the Oracle base directory, you must not use a subdirectory of the Oracle base directory as the CRS home directory.

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 at least 1 MB of free disk space.

If you are using the same file system for the Oracle base directory, this 1 MB of disk space is additional to the free disk space requirement that you identified previously.

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

The path to the CRS home 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 CRS home 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/crs/oracle_sw_owner/product/10.1.0/crs
# chown -R oracle:oinstall /mount_point/crs/oracle_sw_owner
# chmod -R 775 /mount_point/crs/oracle_sw_owner
```

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

```
/u01/crs/oracle/product/10.1.0/crs
```

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 installing Oracle CRS, set the ORACLE_HOME environment variable to specify this directory.

Note: During the installation, before you run the `root.sh` script, you must change the permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories.

Choose a Storage Option for Oracle CRS, Database, and Recovery Files

The following table shows the storage options supported for storing Oracle Cluster Ready Services (CRS) files, Oracle database files, and Oracle database recovery files. Oracle database files include datafiles, control files, redo log files, the server parameter file, and the password file. Oracle CRS files include the Oracle Cluster Registry (OCR) and the CRS voting disk.

For all installations, you must choose the storage option that you want to use for Oracle CRS files and Oracle database files. If you want to enable automated backups during the installation, you must also choose the storage option that you want to use for recovery files (the flash recovery area). You do not have to use the same storage option for each file type.

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

<http://metalink.oracle.com>

Storage Option	File Types Supported		
	CRS	Database	Recovery
Automatic Storage Management	No	Yes	Yes
Cluster File System with TruCluster V5.1B	Yes	Yes	Yes
Logical Storage Manager with TruCluster V5.1B	Voting disk only	Yes	No

Storage Option	File Types Supported		
	CRS	Database	Recovery
Shared raw partitions	Voting disk only	Yes	No

Use the following guidelines when choosing the storage options that you want to use for each file type:

- You can choose any combination of the supported storage options for each file type as long as you satisfy any requirements listed for the chosen storage options.
- The Oracle Cluster Registry must be stored on a cluster file system. A shared Logical Storage Manager volume or shared disk partition is not a supported storage option for the OCR.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For Standard Edition installations, ASM is the only supported storage option for database or recovery files.
- You cannot use Automatic Storage Management to store Oracle CRS files, because these files must be accessible before any Oracle instance starts.

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

- To use a file system for Oracle CRS, database, or recovery file storage, refer to the ["Create Directories for Oracle CRS, Database, or Recovery Files"](#) section on page 4-23.
- To use ASM for database or recovery file storage, refer to the ["Configure Disks for Automatic Storage Management"](#) section on page 4-26.

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 on shared disks, for example, a cluster file system.
- 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 that 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 on shared disks, for example, a cluster file system.
- 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 that 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 that 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 on shared disks, for example, a cluster file system.
- 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, refer to the *Oracle Backup and Recovery Basics* manual.

- The default path suggested by the Installer for the flash recovery area 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 that 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: The mount point that you use for the file system must be identical on each node. Make sure that the file systems are configured to mount automatically when a node restarts.

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
Database files	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 for storage, refer to the ["Configure Disks for Automatic Storage Management"](#) section on page 4-26.

Otherwise, refer to the ["Verify that the Cluster Software is Running"](#) section on page 4-32.

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](#)
- [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. Refer to the *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, recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. You can use the file system for one file type and ASM for the other.

For RAC installations, if you choose to enable automated backups and you do not have a shared 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 on 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 database files 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 method 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 database files 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	Database Files	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

For RAC installations, 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 ASM instance is already running 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 protect against 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, refer to the "[Configuring Disks for ASM](#)" section on page 4-30.

Using an Existing ASM Disk Group

If you want to store either database or recovery files in an existing ASM disk group, 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 method 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 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 restart 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`).

3. From the list, identify the device names for the disk devices that you want to add to a disk group.
4. 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.

5. 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.

6. Enter the following command to verify that the device is not mounted as a UFS file system:

```
# /sbin/mount
```

The device that you want to add to the ASM disk group must not be listed in the output.

7. 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.

8. 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, make sure that you set the permissions only on the correct logical device name for the disk.

If the nodes are configured differently, the device name for a particular device might be different on some nodes. Make sure that you specify the correct device names on each node.

9. Repeat steps 2 to 8 on the other cluster nodes.

Verify that the Cluster Software is Running

To verify that the TruCluster software is running on Tru64 UNIX systems, 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, refer to 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/tnslnsr 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, refer to the ["Configure the oracle User's Environment"](#) section on page 4-33 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 (csh or tcsh):

```
% vi .login
```
7. Enter or edit the following line, specifying a value of 022 for the default file mode 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 that 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 that 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 or identified earlier.

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

■ Bourne, Bash, or Korn shell:

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

■ C shell:

```
% unsetenv ORACLE_HOME
% unsetenv TNS_ADMIN
```

Note: If the ORACLE_HOME environment variable is set, the Installer uses the value that 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.

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

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

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

Pre-Installation Tasks for RAC on Linux

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
- Create the CRS Home Directory
- Choose a Storage Option for Oracle CRS, Database, and Recovery Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Configure Raw Devices
- 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-mode 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 silent-mode installations, refer to the *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 are not installing the software on the local system, use the `ssh`, `rlogin`, or `telnet` command to connect to the system where you want to install the software:

```
$ 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, refer to 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

Each 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 and platform
- 1.2 GB of disk space for a preconfigured database that uses file system storage (optional)

Note: The disk space requirement 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 each 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, refer to 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 -h /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 -h
```

5. To determine whether the system architecture can run the software, enter the following 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, refer to 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 non-routable IP addresses for the private 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	143.47.43.104	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

Refer to 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.

Checking the Software Requirements on Linux

Check that the required software and patches are installed on the system.

Check for Required Software

Depending on the system architecture, verify that the software listed in one of the following tables is installed on the system:

- [Table 5–1, "Software Requirements for Installations on x86 Systems"](#)
- [Table 5–2, "Software Requirements for Installations on Itanium Systems"](#)
- [Table 5–3, "Software Requirements for Installations on zSeries Systems"](#)

Additionally, depending on the products that you intend to install or use, verify that the software listed in [Table 5–4, "Product-Specific Software Requirements for Linux Systems"](#) is installed.

The procedure following the tables describes how to check these requirements.

Table 5–1 Software Requirements for Installations on x86 Systems

Item	Requirement
Operating system	<p>One of the following operating system versions:</p> <ul style="list-style-type: none"> ■ Red Hat Enterprise Linux AS/ES 2.1 (Update 3 or higher) ■ Red Hat Enterprise Linux AS/ES 3 (Update 2 or higher) ■ SUSE Linux Enterprise Server (SLES) 8 (service pack 3 or higher) ■ SUSE Linux Enterprise Server 9
Kernel version	<p>The system must be running the following kernel version (or a higher version):</p> <p>Red Hat Enterprise Linux 2.1: 2.4.9, errata 34 (for example 2.4.9-e.34)</p> <p>Red Hat Enterprise Linux 3: 2.4.21-15.EL</p> <p>SUSE Linux Enterprise Server 8: 2.4.21-251</p> <p>SUSE Linux Enterprise Server 9: 2.6.5-7.5</p>

Table 5–1 (Cont.) Software Requirements for Installations on x86 Systems

Item	Requirement
Packages	<p>The following packages (or higher versions) must be installed:</p> <p>Red Hat Enterprise Linux 2.1:</p> <pre>make-3.79.1 glibc-2.2.4-32.15 gcc-2.96-128 gcc-c++-2.96-128 libstdc++-2.96-128 openmotif-2.1.30-11 shadow-utils-20000902-16</pre> <p>Red Hat Enterprise Linux 3:</p> <pre>make-3.79.1 gcc-3.2.3-34 glibc-2.3.2-95.20 compat-db-4.0.14-5 compat-gcc-7.3-2.96.128 compat-gcc-c++-7.3-2.96.128 compat-libstdc++-7.3-2.96.128 compat-libstdc++-devel-7.3-2.96.128 openmotif21-2.1.30-8 setarch-1.3-1</pre> <p>SUSE Linux Enterprise Server 8:</p> <pre>make-3.79.1 gcc-3.2.2-38 gcc-c++-3.2.2-38 openmotif-2.2.2-124</pre> <p>SUSE Linux Enterprise Server 9:</p> <pre>gcc-3.3.3-43 gcc-c++-3.3.3-43 glibc-2.3.3-98 libaio-0.3.98-18 libaio-devel-0.3.98-18 make-3.80 openmotif-libs-2.2.2-519.1</pre>

Table 5–2 Software Requirements for Installations on Itanium Systems

Item	Requirement
Operating system	<p>One of the following operating system versions:</p> <ul style="list-style-type: none"> ■ Red Hat Enterprise Linux AS/ES 2.1 (Update 3 or higher) ■ Red Hat Enterprise Linux AS/ES 3 (Update 1 or higher) ■ SUSE Linux Enterprise Server (SLES) 8 (service pack 3 or higher) ■ SUSE Linux Enterprise Server 9

Table 5–2 (Cont.) Software Requirements for Installations on Itanium Systems

Item	Requirement
Kernel version	<p>The system must be running the following kernel version (or a higher version):</p> <p>Red Hat Enterprise Linux 2.1: 2.4.18, errata 37 (for example 2.4.18-e.37)</p> <p>Red Hat Enterprise Linux 3: 2.4.21-4.EL</p> <p>SUSE Linux Enterprise Server 8: 2.4.21-241</p> <p>SUSE Linux Enterprise Server 9: 2.6.5-7.5</p>
Packages	<p>The following packages (or higher versions) must be installed:</p> <p>Red Hat Enterprise Linux 2.1: make-3.79.1 glibc-2.2.4-32 gcc-2.96-128 gcc-c++-2.96-128 libstdc++-2.96-128 openmotif-2.1.30-11 libaio-0.3.92-1 libaio-devel-0.3.92-1</p> <p>Red Hat Enterprise Linux 3: make-3.79.1 gcc-3.2.3-20 gcc-c++-3.2.3-20 glibc-2.3.2-95.3 compat-db-4.0.14-5 compat-gcc-7.3-2.96.128 compat-gcc-c++-7.3-2.96.128 compat-libstdc++-7.3-2.96.128 compat-libstdc++-devel-7.3-2.96.128 openmotif21-2.1.30-8 setarch-1.3-1</p> <p>SUSE Linux Enterprise Server 8: make-3.79.1 gcc-3.2.2-23 gcc-c++-3.2.2-23 glibc-2.2.5-161 openmotif-2.2.2-125</p> <p>SUSE Linux Enterprise Server 9: gcc-3.3.3-43 gcc-c++-3.3.3-43 glibc-2.3.3-98 libaio-0.3.98-18 libaio-devel-0.3.98-18 make-3.80 openmotif-libs-2.2.2-519.1</p>

Table 5–3 Software Requirements for Installations on zSeries Systems

Item	Requirement
Operating system	The following operating system version: <ul style="list-style-type: none"> ■ SUSE Linux Enterprise Server (SLES) 8 (service pack 3 or higher)
Kernel version	The system must be running the following kernel version (or a higher version): SUSE Linux Enterprise Server 8: 2.4.21-112
Packages	The following packages (or higher versions) must be installed: SUSE Linux Enterprise Server 8: gcc-3.2.2-29 gcc-c++-3.2.2-29 glibc-2.2.5-108 glibc-devel-32bit-8.1-9 libaio-0.3.15-106 make-3.79.1 binutils-2.14.90.0.5-41 openmotif-2.2.2-52

Table 5–4 Product-Specific Software Requirements for Linux Systems

Installation Type or Product	Requirement
Oracle Spatial	X Window System development package: <ul style="list-style-type: none"> ■ Red Hat and SLES 9: XFree86 XFree86-devel ■ SLES 8: xf86 x-devel <p>Note: This software is required only to build the sample programs.</p>
Oracle Messaging Gateway (x86 and zSeries systems only)	Oracle Messaging Gateway supports the integration of Oracle Streams Advanced Queuing (AQ) with the following software: <ul style="list-style-type: none"> ■ IBM WebSphere MQ V5.3, client and server, with corrective service diskette 5 (CSD05) or higher: MQSeriesClient MQSeriesServer MQSeriesRuntime ■ TIBCO Rendezvous 7.2 (x86 only)
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	The version of GNU C and C++ compilers listed previously for your distribution is supported for use with these products. Notes: Intel C++ Compiler v7.1.0.28 or higher is also supported, however it is not required for installation. On Red Hat Enterprise Linux 3, OCCI is supported with both version 2.96 and version 3.2 of the GNU C++ compiler, however, version 3.2 is the default compiler version.

Table 5–4 (Cont.) Product-Specific Software Requirements for Linux Systems

Installation Type or Product	Requirement
Oracle JDBC/OCI Drivers	<p>On x86 systems only, you can use the following optional JDK version with the Oracle JDBC/OCI drivers, however it is not required for the installation:</p> <ul style="list-style-type: none"> ■ Sun JDK 1.3.1_09 with the JNDI extension <p>Note: On x86 and Itanium systems, JDK 1.4.2_04 is installed with this release. On zSeries systems, JDK 1.4.2 (31-bit) is installed.</p>
Real Application Clusters	<p>Oracle Cluster File System (OCFS):</p> <ul style="list-style-type: none"> ■ Version 1.0.13-1 or higher for SUSE Linux Enterprise Server 8 ■ Version 1.0.11-1 or higher for other installations <p>Note: OCFS is required only if you want to use a cluster file system for database file storage. If you want to use Automatic Storage Management or raw devices for database file storage, you do not need to install OCFS.</p> <p><code>ocfs-support</code> <code>ocfs-tools</code> <code>ocfs-kernel_version</code></p>

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. To determine whether the required kernel version is installed, enter the following command:

```
# uname -r
```

If the kernel version is lower than the required version, download and install the required version or a higher version from your Linux vendor's Web site.

3. To determine whether the required packages are installed, enter commands similar to the following:

```
# rpm -q package_name
```

If a package is not installed, install it from your Linux distribution media or download the required package version from your Linux vendor's Web site.

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/>

5. If you require a CSD for WebSphere MQ, refer to 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 refer to 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:

- x86 or Itanium systems:

```
# more /etc/oraInst.loc
```
- zSeries systems:

```
# 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, enter the following command to create the `oinstall` group:

```
# /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 OSDBA 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, enter the following command to create the `dba` group. Use 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, enter the following command to create the `oper` group. Use 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. Refer to 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, refer to the [Verifying that the UNIX User nobody Exists](#) section on page 5-14.
- To modify an existing user, refer to the "Modify an Existing Oracle Software Owner User" section on page 5-14.
- To create a new user, refer to 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. Use 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
```

Refer to the "[Verifying that the UNIX User nobody Exists](#)" section on page 5-14 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, enter a command similar to the following to modify it. Specify 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 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 run the Installer, 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.

If you see any other messages or text, apart from the date, the installation might fail. Make any changes required to ensure that only the date is displayed when you enter these commands.

5. To ensure that X11 forwarding will not cause the installation to fail, create a user-level SSH client configuration file for the Oracle software owner user, as follows:
 - a. Using any text editor, edit or create the `~oracle/.ssh/config` file.
 - b. Make sure that the `ForwardX11` attribute is set to `no`, for example:

```
Host *
    ForwardX11 no
```

6. 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

Refer to the following section for information about configuring the system kernel parameters and shell limits.

Configuring Kernel Parameters on Linux

Note: 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. Refer to your operating system documentation for more information about tuning kernel parameters.

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	Value	File
semmsl	250	/proc/sys/kernel/sem
semmns	32000	
semopm	100	
semmni	128	
shmall	2097152	/proc/sys/kernel/shmall
shmmax	Half the size of physical memory (in bytes)	/proc/sys/kernel/shmmax
shmmni	4096	/proc/sys/kernel/shmmni
file-max	65536	/proc/sys/fs/file-max
ip_local_port_range	1024 65000	/proc/sys/net/ipv4/ip_local_port_range
rmem_default	262144	/proc/sys/net/core/rmem_default
rmem_max	262144	/proc/sys/net/core/rmem_max
wmem_default	262144	/proc/sys/net/core/wmem_default
wmem_max	262144	/proc/sys/net/core/wmem_max

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 table 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.
rmem_default, rmem_max, wmem_default, and wmem_max	# /sbin/sysctl -a grep net.core

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
net.core.rmem_default = 262144
net.core.rmem_max = 262144
net.core.wmem_default = 262144
net.core.wmem_max = 262144
```

By specifying the values in the `/etc/sysctl.conf` file, they persist when you restart 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 SUSE systems only, enter the following command to cause the system to read the `/etc/sysctl.conf` file when it restarts:

```
# /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 on Linux systems, you must increase the following shell limits for the `oracle` user:

Shell Limit	Item in limits.conf	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:

```
oracle          soft  nproc  2047
oracle          hard  nproc  16384
oracle          soft  nofile 1024
oracle          hard  nofile 65536
```

2. Add or edit the following line in 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 SUSE 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 shell (`csh` or `tcsh`), add the following lines to the `/etc/csh.login` file (or the `/etc/csh.login.local` file on SUSE 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
```

If the Oracle base directory is on a cluster file system or an NFS file system on a NAS device, you must specify a path for the Oracle Inventory directory on a local file system. The Oracle base directory must be 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. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
/u01/crs/oracle/product/10.1.0/crs
```

Note: Because you must change the permissions of all of the parent directories of the CRS home directory after installing the software to grant write access only to the `root` user, the CRS home directory must not be a subdirectory of the Oracle base 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:

- x86 or Itanium systems:


```
# more /etc/oraInst.loc
```
- zSeries systems:


```
# 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 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 3 GB, depending on the platform
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 4 GB, depending on the platform

To determine the free disk space on the file system where the Oracle base directory is located, enter the following command:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, refer to the "[Create the CRS Home Directory](#)" section on page 5-26.

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, refer to 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 3 GB, depending on the platform
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 4 GB, depending on the platform

To identify an appropriate file system, follow these steps:

1. Use the `df -h` 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/oracle_sw_owner
# chmod -R 775 /mount_point/app/oracle_sw_owner
```

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.

Create the CRS Home Directory

You must create a CRS home directory before installing Oracle CRS. You can create the CRS home directory on the same file system as the Oracle base directory, or you can choose a different file system for this directory. If you choose the same file system as the Oracle base directory, you must not use a subdirectory of the Oracle base directory as the CRS home directory.

To identify an appropriate file system, follow these steps:

1. Use the `df -h` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has at least 1 MB of free disk space.

If you are using the same file system for the Oracle base directory, this 1 MB of disk space is additional to the free disk space requirement that you identified previously.

Note: The file system can be a local file system or an NFS file system on a certified NAS device. Do not create the CRS home directory on an OCFS file system.

The path to the CRS home 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 CRS home 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/crs/oracle_sw_owner/product/10.1.0/crs
# chown -R oracle:oinstall /mount_point/crs/oracle_sw_owner
# chmod -R 775 /mount_point/crs/oracle_sw_owner
```

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

```
/u01/crs/oracle/product/10.1.0/crs
```

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 installing Oracle CRS, set the `ORACLE_HOME` environment variable to specify this directory.

Note: During the installation, before you run the `root.sh` script, you must change the permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories.

Choose a Storage Option for Oracle CRS, Database, and Recovery Files

The following table shows the storage options supported for storing Oracle Cluster Ready Services (CRS) files, Oracle database files, and Oracle database recovery files. Oracle database files include datafiles, control files, redo log files, the server parameter

file, and the password file. Oracle CRS files include the Oracle Cluster Registry (OCR) and the CRS voting disk.

For all installations, you must choose the storage option that you want to use for Oracle CRS files and Oracle database files. If you want to enable automated backups during the installation, you must also choose the storage option that you want to use for recovery files (the flash recovery area). You do not have to use the same storage option for each file type.

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>

Storage Option	File Types Supported		
	CRS	Database	Recovery
Automatic Storage Management	No	Yes	Yes
Cluster file system (OCFS)	Yes	Yes	Yes
NFS file system	Yes	Yes	Yes
Note: Requires a certified NAS device			
Shared raw partitions	Yes	Yes	No
Shared raw logical volumes (zSeries only)	Yes	Yes	No

Use the following guidelines when choosing the storage options that you want to use for each file type:

- You can choose any combination of the supported storage options for each file type provided that you satisfy any requirements listed for the chosen storage options.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For Standard Edition installations, ASM is the only supported storage option for database or recovery files.
- You cannot use Automatic Storage Management to store Oracle CRS files, because these files must be accessible before any Oracle instance starts.

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

- To use a file system for Oracle CRS, database, or recovery file storage, refer to the "[Create Directories for Oracle CRS, Database, or Recovery Files](#)" section on page 5-28.
- To use ASM for database or recovery file storage, refer to the "[Configure Disks for Automatic Storage Management](#)" section on page 5-31.
- To use raw devices (partitions) for Oracle CRS or database file storage, refer to the "[Configure Raw Devices](#)" section on page 5-44.

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 on shared disks, for example:
 - A cluster file system such as OCFS
 - An NFS file system on a certified NAS device

Note: If you are using a shared file system on a NAS device to store a shared Oracle home directory for CRS or RAC, you must use the same NAS device for Oracle CRS file storage.

- 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 that 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 on shared disks, for example:
 - A cluster file system such as OCFS
 - An NFS file system on a certified NAS device

Note: If you are using a shared file system on a NAS device to store a shared Oracle home directory for CRS or RAC, you must use the same NAS device for Oracle database file storage.

- 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 that 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 that 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 on shared disks, for example:
 - A cluster file system such as OCFS
 - 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, refer to the *Oracle Backup and Recovery Basics* manual.

- The default path suggested by the Installer for the flash recovery area 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 that 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: The mount point that you use for the file system must be identical on each node. Make sure that the file systems are configured to mount automatically when a node restarts.

2. Use the `df -h` 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
Database files	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, refer to one of the following sections:

- [Configure Disks for Automatic Storage Management](#) on page 5-31
- [Configure Raw Devices](#) on page 5-44

Otherwise, refer to the "[Verify that the Required Software is Running](#)" section on page 5-53.

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](#)
- [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. Refer to the *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, recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. You can use the file system for one file type and ASM for the other.

For RAC installations, if you choose to enable automated backups and you do not have a shared 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 on 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 database files 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 method 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 database files 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	Database Files	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

For RAC installations, 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 ASM instance is already running 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 protect against 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, refer to the "[Configuring Disks for ASM](#)" section on page 5-35.

Using an Existing ASM Disk Group

If you want to store either database or recovery files in an existing ASM disk group, 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 method 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 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

The following sections describe how to configure the disks that you want to use in an ASM disk group. It includes information about the following:

- [Configuring Disks for ASM on Linux](#) on page 5-35

Configuring Disks for ASM on Linux

On Linux systems, Oracle provides an ASM library driver that you can use to simplify the configuration and management of the disk devices that you want to use with ASM. If you intend to use ASM for database storage on Linux, 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 each method:

- [Configuring Disks for ASM Using the ASM Library Driver](#)
- [Configuring Disks for ASM Using Raw Devices](#)

Note: If you choose to configure disks using the ASM library driver, you must use Database Configuration Assistant (DBCA) in interactive mode to create the database. You can run DBCA in interactive mode by choosing the Custom installation type or the Advanced database configuration option. You must also change the default disk discovery string to `ORCL: *`.

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 software, follow these steps:

1. Enter the following command to determine the kernel version and architecture of the system:

```
# uname -rm
```

2. If necessary, download the required ASM library driver packages from the OTN Web site:

```
http://www.oracle.com/technology/software/tech/linux/asmlib/index.html
```

Note: For release 10.1.0.3, ASM library driver packages for some kernel versions are available on the Oracle Cluster Ready Services CD-ROM in the `RPMS/asmlib` directory and on the DVD-ROM in the `crs/RPMS/asmlib` directory. However, Oracle recommends that you check the OTN Web site for the most up-to-date packages.

You must install the following three packages, where *version* is the version of the ASM library driver, *arch* is the system architecture, and *kernel* is the version of the kernel that you are using:

```
oracleasm-support-version.arch.rpm  
oracleasm-kernel-version.arch.rpm  
oracleasmlib-version.arch.rpm
```

3. Switch user to the `root` user:

```
$ su -
```

4. Enter a command similar to the following to install the packages:

```
# rpm -Uvh oracleasm-support-version.arch.rpm \  
          oracleasm-kernel-version.arch.rpm \  
          oracleasmlib-version.arch.rpm
```

For example, if you are using the Red Hat Enterprise Linux AS 2.1 enterprise kernel on an x86 system, enter a command similar to the following:

```
# rpm -Uvh oracleasm-support-1.0.0-1.i386.rpm \  
          oracleasm-2.4.9-e-enterprise-1.0.0-1.i686.rpm \  
          oracleasmlib-1.0.0-1.i386.rpm
```

5. 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
```

6. Enter the following command to run the `oracleasm` initialization script with the `configure` option:

```
# /etc/init.d/oracleasm configure
```

7. 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.

8. 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, follow these steps:

1. If you intend to use IDE, SCSI, or RAID devices in the ASM disk group, follow these steps:
 - a. If necessary, install or configure the shared disk devices that you intend to use for the disk group and restart the system.
 - b. 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	<code>/dev/hdxn</code>	In this example, <code>x</code> is a letter that identifies the IDE disk and <code>n</code> 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, <code>x</code> is a letter that identifies the SCSI disk and <code>n</code> is the partition number. For example, <code>/dev/sda</code> is the first disk on the first SCSI bus.

Disk Type	Device Name Format	Description
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.

To include devices in a disk group, you can specify either whole-drive device names or partition device names.

Note: On Linux systems, Oracle recommends that you create a single whole-disk partition on each disk that you want to use.

- c. Use either `fdisk` or `parted` to create a single whole-disk partition on the disk devices that you want to use.
2. On zSeries Linux only, if you intend to use direct access storage devices (DASDs), follow these steps:
 - a. If necessary, install or configure the shared DASDs that you intend to use for the disk group and restart the system.
 - b. Enter the following command to identify the DASDs configured on the system:

```
# more /proc/dasd/devices
```

The output from this command contains lines similar to the following:

```
0302(ECKD) at ( 94: 48) is dasdm : active at blocksize: 4096, 540000 blocks, 2109 MB
```

These lines display the following information for each DASD:

- The device number (0302)
- The device type (ECKD or FBA)
- The Linux device major and minor numbers (94: 48)
- The Linux device file name (dasdm)

In general, DASDs have device names in the form `dasdxxxx`, where `xxxx` is between one and four letters that identify the device.

- The block size and size of the device

- c. From the display, identify the devices that you want to use.
- d. If you want to use FBA type DASDs, go to step 3.

You do not have to format FBA type DASDs in Linux. The device name for the single whole-disk partition for FBA type DASDs is `/dev/dasdxxxx1`.

- e. If you want to use ECKD type DASDs, enter a command similar to the following to format the DASD, if it is not already formatted:

```
# /sbin/dasdfmt -b 4096 -f /dev/dasdxxxx
```

Caution: Formatting a DASD destroys all existing data on the device. Make sure that:

- You specify the correct DASD device name
 - The DASD does not contain existing data that you want to preserve
-
-

This command formats the DASD with a block size of 4 KB and the compatible disk layout (default), which enables you to create up to three partitions on the DASD.

Alternatively, you could use the `-d 1d1` option to format the DASD using the Linux disk layout if you require only a single partition. If you use this disk layout, the partition device name for the DASD is `/dev/dasdxxxx1`.

- f. If you formatted the DASD with the compatible disk layout, enter a command similar to the following to create a single whole-disk partition on the device:

```
# /sbin/fdasd -a /dev/dasdxxxx
```

The device name for the single whole-disk partition for the DASDs is `/dev/dasdxxxx1`

3. 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: The disk names that you specify can contain uppercase letters, numbers, and the underscore character. They must start with an uppercase letter.

If you are using a multi-pathing disk driver with ASM, make sure that you specify the correct logical device name for the disk.

4. 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.

Note: To create a database during the installation using the ASM library driver, you must choose an installation method that runs DBCA in interactive mode, for example by choosing the Custom installation type or the Advanced database configuration option. You must also change the default disk discovery string to `ORCL:*`.

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: # /etc/init.d/oracleasm configure
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: # /etc/init.d/oracleasm enable
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 restarting the system: # /etc/init.d/oracleasm restart
createdisk	Use the <code>createdisk</code> option to mark a disk device for use with the ASM library driver and give it a name: # /etc/init.d/oracleasm createdisk <i>DISKNAME</i> <i>devicename</i>
deletedisk	Use the <code>deletedisk</code> option to unmark a named disk device: # /etc/init.d/oracleasm deletedisk <i>DISKNAME</i> 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: # /etc/init.d/oracleasm querydisk { <i>DISKNAME</i> <i>devicename</i> }
listdisks	Use the <code>listdisks</code> option to list the disk names of marked ASM library driver disks: # /etc/init.d/oracleasm listdisks
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: # /etc/init.d/oracleasm scandisks

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:

1. If you intend to use IDE, SCSI, or RAID devices in the ASM disk group, follow these steps:
 - a. If necessary, install or configure the shared disk devices that you intend to use for the disk group and restart the system.
 - b. 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	<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.

To include devices in a disk group, you can specify either whole-drive device names or partition device names.

Note: On Linux systems, Oracle recommends that you create a single whole-disk partition on each disk that you want to use.

- c. Use either `fdisk` or `parted` to create a single whole-disk partition on the disk devices that you want to use.
2. On zSeries Linux only, if you intend to use direct access storage devices (DASDs), follow these steps:
 - a. If necessary, install or configure the shared DASDs that you intend to use for the disk group and restart the system.
 - b. Enter the following command to identify the DASDs configured on the system:

```
# more /proc/dasd/devices
```

The output from this command contains lines similar to the following:

```
0302(ECKD) at ( 94: 48) is dasdm : active at blocksize: 4096, 540000 blocks, 2109 MB
```

These lines display the following information for each DASD:

- The device number (0302)
- The device type (ECKD or FBA)
- The Linux device major and minor numbers (94: 48)
- The Linux device file name (dasdm)

In general, DASDs have device names in the form `dasdxxxx`, where `xxxx` is between one and four letters that identify the device.

- The block size and size of the device

- c. From the display, identify the devices that you want to use.
- d. If you want to use FBA type DASDs, go to step 3.

You do not have to format FBA type DASDs in Linux. The device name for the single whole-disk partition for FBA type DASDs is `/dev/dasdxxxx1`.

- e. If you want to use ECKD type DASDs, enter a command similar to the following to format the DASD, if it is not already formatted:

```
# /sbin/dasdfmt -b 4096 -f /dev/dasdxxxx
```

Caution: Formatting a DASD destroys all existing data on the device. Make sure that:

- You specify the correct DASD device name
 - The DASD does not contain existing data that you want to preserve
-
-

This command formats the DASD with a block size of 4 KB and the compatible disk layout (default), which enables you to create up to three partitions on the DASD.

Alternatively, you could use the `-d 1d1` option to format the DASD using the Linux disk layout if you require only a single partition. If you use this disk layout, the partition device name for the DASD is `/dev/dasdxxxx1`.

- f. If you formatted the DASD with the compatible disk layout, enter a command similar to the following to create a single whole-disk partition on the device:

```
# /sbin/fdasd -a /dev/dasdxxxx
```

The device name for the single whole-disk partition for the DASDs is `/dev/dasdxxxx1`

3. On Red Hat systems, complete the following steps on each node to bind the disk devices to raw devices:

Note: If the nodes are configured differently, the disk device names might be different on some nodes. In the following procedure, make sure that you specify the correct disk device names on each node.

- a. To determine what raw devices are already bound to other devices, enter the following command on every node:

```
# /usr/bin/raw -qa
```

Raw devices have device names in the form `/dev/raw/rawn`, where *n* is a number that identifies the raw device.

For each device that you want to include in the disk group, identify a raw device name that is unused on all nodes.

- 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, make sure 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 `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 disk devices to the raw devices, enter the following command:

```
# /sbin/service rawdevices restart
```

The system automatically binds the devices listed in the `rawdevices` file when it restarts.

4. On SUSE systems, complete the following steps on each node to bind the disk devices to raw devices:

Note: If the nodes are configured differently, the disk device names might be different on some nodes. In the following procedure, make sure that you specify the correct disk device names on each node.

- a. To determine what raw devices are already bound to other devices, enter the following command on every node:

```
# /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.

For each device that you want to include in the disk group, identify a raw device name that is unused on all nodes.

- 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, make sure 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 start
```

- e. To ensure that the raw devices are bound when the system restarts, enter the following command:

```
# /sbin/chkconfig raw on
```

- 5. If you also want to use raw devices for storage, refer to the ["Configure Raw Devices"](#) section on page 5-44.

Otherwise, refer to the ["Verify that the Required Software is Running"](#) section on page 5-53.

Configure Raw Devices

Refer to the following section for information about configuring raw devices.

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

Configuring Raw Partitions or Raw Logical Volumes on Linux

The following subsections describe how to configure raw partitions or raw logical volumes on Linux.

Review Important Information

The procedures in the following sections describe how to create raw partitions or raw logical volumes for Oracle CRS and database file storage. Raw logical volumes are supported for RAC only on zSeries systems. On x86 and Itanium systems, Oracle does not support the use of logical volumes with RAC for either CRS or database files.

Identify the Required Raw Partitions

[Table 5-5](#) lists the number and size of the raw partitions or logical volumes that you must configure for database files. [Table 5-6](#) lists the number and size of the raw partitions or logical volumes that you must configure for CRS files.

Note: On x86 and Itanium systems, the use of logical volumes for raw devices is supported only for single-instance databases. Their use is not supported for RAC databases on these platforms.

Table 5-5 Raw Partitions or Logical Volumes Required for Database Files

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	500	SYSTEM tablespace: <i>dbname_system_raw_500m</i>
1	300 + (Number of instances * 250)	SYSAUX tablespace: <i>dbname_sysaux_raw_800m</i>

Table 5–5 (Cont.) Raw Partitions or Logical Volumes Required for Database Files

Number	Size (MB)	Purpose and Sample Logical Volume Name
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	250	TEMP tablespace: <i>dbname_temp_raw_250m</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_redoln_m_raw_120m</i>
2	110	First and second control files: <i>dbname_control{1 2}_raw_110m</i>
1	5	Server parameter file (SPFILE): <i>dbname_spfile_raw_5m</i>
1	5	Password file: <i>dbname_pwdfile_raw_5m</i>

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.

Table 5–6 Raw Partitions or Logical Volumes Required for CRS Files

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	100	Oracle Cluster Registry <i>ora_ocr_raw_100m</i> 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 (OCR). 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 <i>ora_vote_raw_20m</i> 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.

To create the required raw partitions or logical volumes, refer to one of the following sections:

- [Creating the Required Raw Partitions on IDE, SCSI, or RAID Devices](#) on page 5-46

- [zSeries Linux Only: Creating the Required Raw Logical Volumes](#) on page 5-47
- [zSeries Linux Only: Creating the Required Raw Partitions on DASDs](#) on page 5-49

Creating the Required Raw Partitions on IDE, SCSI, or RAID Devices

If you intend to use IDE, SCSI, or RAID devices for the raw devices, follow these steps:

1. If necessary, install or configure the shared disk devices that you intend to use for the raw partitions and restart 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, 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.

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.

3. 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 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.

4. To continue, refer to the ["Bind the Partitions or Logical Volumes to Raw Devices"](#) section on page 5-51.

zSeries Linux Only: Creating the Required Raw Logical Volumes

On zSeries Linux only, you can use raw logical volume manager (LVM) volumes for Oracle CRS and database file storage. You can create the required raw logical volumes in a volume group on either direct access storage devices (DASDs) or on SCSI devices. To configure the required raw logical volumes, follow these steps:

1. If you intend to use DASDs for the volume group, follow these steps:
 - a. If necessary, install or configure the shared DASDs that you intend to use for the volume group and restart the system.
 - b. Enter the following command to identify the DASDs configured on the system:

```
# more /proc/dasd/devices
```

The output from this command contains lines similar to the following:

```
0302(ECKD) at ( 94: 48) is dasdm : active at blocksize: 4096, 540000 blocks, 2109 MB
```

These lines display the following information for each DASD:

- The device number (0302)
- The device type (ECKD or FBA)
- The Linux device major and minor numbers (94: 48)
- The Linux device file name (dasdm)

In general, DASDs have device names in the form `dasdxxxx`, where `xxxx` is between one and four letters that identify the device.

- The block size and size of the device

- c. From the display, identify the devices that you want to use.
- d. If you want to use ECKD type DASDs, enter a command similar to the following to format the DASD, if it is not already formatted:

```
# /sbin/dasdfmt -b 4096 -f /dev/dasdxxxx
```

Caution: Formatting a DASD destroys all existing data on the device. Make sure that:

- You specify the correct DASD device name
 - The DASD does not contain existing data that you want to preserve
-

This command formats the DASD with a block size of 4 KB and the compatible disk layout (default), which enables you to create up to three partitions on the DASD.

Note: For the DASDs that you intend to use to store the Oracle CRS files (the Oracle Cluster Registry and the CRS voting disk) you must use a 4 KB block size.

Alternatively, you could use the `-d ld1` option to format the DASD using the Linux disk layout.

- e. If you formatted the DASD with the compatible disk layout, enter a command similar to the following to create a single whole-disk partition on the device:

```
# /sbin/fdasd -a /dev/dasdxxxx
```

The device name for the single whole-disk partition for the DASDs is `/dev/dasdxxxx1`.

- f. To continue, go to step 3.

Note: The examples in the following steps show SCSI device names. You must use the appropriate DASD device names when completing these steps.

2. If you intend to use SCSI devices in the volume group, follow these steps:
 - a. If necessary, install or configure the shared disk devices that you intend to use for the volume group and restart the system.
 - b. To identify the device name for the disks that you want to use, enter the following command:

```
# /sbin/fdisk -l
```

SCSI devices have device names similar to the following:

```
/dev/sdxn
```

In this example, *x* is a letter that identifies the SCSI disk and *n* is the partition number. For example, `/dev/sda` is the first disk on the first SCSI bus.

- c. If necessary, use `fdisk` to create partitions on the devices that you want to use.
 - d. Use the `t` command in `fdisk` to change the system ID for the partitions that you want to use to `0x8e`.
3. Enter a command similar to the following to mark each device that you want to use in the volume group as a physical volume:

```
# pvcreate /dev/sda1 /dev/sdb1
```

4. To create a volume group named `oracle_vg` using the devices that you marked, enter a command similar to the following:

```
# vgcreate oracle_vg /dev/dasda1 /dev/dasdb1
```

5. To create the required logical volumes in the volume group that you created, enter commands similar to the following:

```
# lvcreate -L size -n lv_name vg_name
```

In this example:

- *size* is the size of the logical volume, for example `500M`
- *lv_name* is the name of the logical volume, for example `orcl_system_raw_500m`
- *vg_name* is the name of the volume group, for example `oracle_vg`

For example, to create a 500 MB logical volume for the SYSTEM tablespace for a database named rac in the oracle_vg volume group, enter the following command:

```
# lvcreate -L 500M -n rac_system_raw_500m oracle_vg
```

Note: These commands create a device name similar to the following for each logical volume:

```
/dev/vg_name/lv_name
```

6. On the other cluster nodes, enter the following commands to configure the volume group and logical volumes on those nodes:

```
# vgscan
# vgchange -a y
```

7. To continue, refer to the "[Bind the Partitions or Logical Volumes to Raw Devices](#)" section on page 5-51.

zSeries Linux Only: Creating the Required Raw Partitions on DASDs

On zSeries Linux only, if you intend to use direct access storage devices (DASDs) for the raw devices, follow these steps:

1. If necessary, install or configure the shared DASDs that you intend to use for the raw partitions and restart the system.

Note: Not all DASD types support more than one partition. Also, for DASDs that do support more than one partition, the maximum number of partitions supported is three.

If possible, configure DASDs with sizes that match the raw devices that you want to use to avoid wasting disk space.

2. Enter the following command to identify the DASDs configured on the system:

```
# more /proc/dasd/devices
```

The output from this command contains lines similar to the following:

```
0302(ECKD) at ( 94: 48) is dasdm : active at blocksize: 4096, 540000 blocks, 2109 MB
```

These lines display the following information for each DASD:

- The device number (0302)
- The device type (ECKD or FBA)
- The Linux device major and minor numbers (94: 48)
- The Linux device file name (dasdm)

In general, DASDs have device names in the form dasdxxxx, where xxxx is between one and four letters that identify the device.

- The block size and size of the device

3. From the display, identify the devices that you want to use.

4. If you want to use FBA type DASDs, go to the ["Bind the Partitions or Logical Volumes to Raw Devices"](#) section on page 5-51.

You do not have to format FBA type DASDs in Linux. The device name for the single whole-disk partition for FBA type DASDs is `/dev/dasdxxxx1`.

5. If you want to use ECKD type DASDs, enter a command similar to the following to format the DASD, if it is not already formatted:

```
# /sbin/dasdfmt -b 4096 -f /dev/dasdxxxx
```

Caution: Formatting a DASD destroys all existing data on the device. Make sure that:

- You specify the correct DASD device name
 - The DASD does not contain existing data that you want to preserve
-
-

This command formats the DASD with a block size of 4 KB and the compatible disk layout (default), which enables you to create up to three partitions on the DASD.

Note: For the DASDs that you intend to use to store the Oracle CRS files (the Oracle Cluster Registry and the CRS voting disk) you must use a 4 KB block size.

Alternatively, you could use the `-d 1d1` option to format the DASD using the Linux disk layout if you require only a single partition. If you use this disk layout, the partition device name for the DASD is `/dev/dasdxxxx1`.

6. If you formatted the DASD with the compatible disk layout, enter a command similar to the following to create up to three partitions on the device, if required:

```
# /sbin/fdasd /dev/dasdxxxx
```

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 the required partitions on this device, use the `w` command to write the modified partition table to the device.
- See the `fdasd` man page for more information about creating partitions.

The partitions on a DASD have device names similar to the following, where *n* is the partition number, between 1 and 3:

```
/dev/dasdxxxxn
```

Note: The examples in the following sections show SCSI device names. You must use the appropriate DASD device names when completing these procedures.

Bind the Partitions or Logical Volumes to Raw Devices

After you have created the required partitions or logical volumes, you must bind them to raw devices on every node. 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 Linux distribution that you are using:

Note: If the nodes are configured differently, the disk device names might be different on some nodes. In the following procedure, make sure that you specify the correct disk device names on each node.

- Red Hat:

1. To determine what raw devices are already bound to other devices, enter the following command on every node:

```
# /usr/bin/raw -qa
```

Raw devices have device names in the form `/dev/raw/rawn`, where `n` is a number that identifies the raw device.

For each device that you want to use, identify a raw device name that is unused on all nodes.

2. Open the `/etc/sysconfig/rawdevices` file in any text editor and add a line similar to the following for each partition or logical volume that you created:

```
/dev/raw/raw1 /dev/sdb1
```

Specify an unused raw device for each partition or logical volume.

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:oinstall /dev/raw/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:dba /dev/raw/rawn
# chmod 660 /dev/raw/rawn
```

5. To bind the partitions or logical volumes to the raw devices, enter the following command:

```
# /sbin/service rawdevices restart
```

The system automatically binds the devices listed in the `rawdevices` file when it restarts.

6. Repeat steps 2 to 5 on the other cluster nodes.

- SUSE:

1. To determine what raw devices are already bound to other devices, enter the following command on every node:

```
# /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.

For each device that you want to use, identify a raw device name that is unused on all nodes.

2. Open the `/etc/raw` file in any text editor and add a line similar to the following to associate each partition or logical volume 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:oinstall /dev/raw/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 or logical volumes to the raw devices, enter the following command:

```
# /etc/init.d/raw start
```

6. To ensure that the raw devices are bound when the system restarts, enter the following command:

```
# /sbin/chkconfig raw on
```

7. Repeat steps 2 to 6 on the other cluster nodes.

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. 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 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 tablespace datafiles.
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 on Linux systems, 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.

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 restarts 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 restarts the system. Using the default values, the node would be restarted 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 restarts, 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.
 - SUSE:
On SUSE 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, refer to 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.

Note: If no Oracle Net listeners are running, refer to the "[Configure the oracle User's Environment](#)" section on page 5-56 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 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:

- Bash shell on Red Hat Enterprise Linux:

```
$ vi .bash_profile
```

- 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 mode 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:

- Bash shell on Red Hat Enterprise 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 that 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 -h` command to identify a suitable file system with sufficient free space.

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

```
$ 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 or identified earlier.

14. If you are using raw devices for database file 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:

- Bourne, Bash, or Korn shell:

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

- C shell:

```
% unsetenv ORACLE_HOME
% unsetenv TNS_ADMIN
```

Note: If the `ORACLE_HOME` environment variable is set, the Installer uses the value that 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 that the environment variables you set in this section have the correct values.

Pre-installation Tasks for RAC on Mac OS X

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](#)
- [Create the CRS Home Directory](#)
- [Choose a Storage Option for Oracle CRS, Database, and Recovery Files](#)
- [Configure Disks for Automatic Storage Management](#)
- [Configure Raw Partitions](#)
- [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, follow these steps:

1. Log in to the system as a user that is a member of the `admin` group.
2. Start a local terminal session (`/Applications/Utilities/Terminal`).
3. If you are not logged in as the `root` user, enter the following command to start a shell session with `root` privileges:

```
$ sudo sh
```
4. When prompted, enter your password.

Note: Alternatively, if you are logged in as a user who is a member of the `admin` group, you can use the `sudo` command to run commands as the `root` user, as follows:

```
$ sudo /usr/sbin/sysctl -a
```

Check the Hardware Requirements

Each system must meet the following minimum hardware requirements:

- 512 MB of physical RAM (536870912 bytes)
- 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 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 requirement 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 each system meets these requirements, follow these steps:

1. To determine the physical RAM size, use System Profiler (`/Applications/Utilities/System Profiler`) or enter the following command:

```
# /usr/sbin/system_profiler SPHardwareDataType | grep 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 ensure that the system has enough swap space, enter the following command:

```
# df -h /
```

Note: Mac OS X dynamically creates swap files as required in the `/private/var/vm` directory. Make sure that you have at least 1 GB of available disk space on the root (`/`) file system to accommodate newly created swap files.

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

```
# df -h /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 -h
```

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 non-routable IP addresses for the private interfaces, for example 10.*.* or 192.168.*.*. You can use the `/etc/hosts` file on each node (or NetInfo) 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	143.47.43.104	DNS
<code>rac1-priv1</code>	Private	10.0.0.1	<code>/etc/hosts</code> or NetInfo
<code>rac1-priv2</code>	Private	10.0.0.2	<code>/etc/hosts</code> or NetInfo

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, complete one of the following steps:
 - 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
```
 - Enter commands similar to the following on all nodes to register the private IP address and private host name in the local NetInfo database:


```
# nictl . create /machines/rac1-priv1
# nictl . append /machines/rac1-priv1 ip_address 10.0.0.1
```
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

Check that the required software and patches are installed on the system.

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	The following operating system version: Apple Mac OS X Server 10.3.6
All Installations	The following software (or a higher version) must be installed: <ul style="list-style-type: none"> ■ Xcode Tools 1.1 or higher ■ GNU C compiler (<code>gcc</code>) version 3.3 20030304 (Apple Computer, Inc. build 1635) <p>This version of <code>gcc</code> is included in Xcode Tools 1.2 or higher or in the Dec 2003 <code>gccLongBranch</code> Tools update.</p>
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	The version of the GNU C and C++ compiler listed previously is supported for use with these products.
Oracle Spatial	The following software is required to build the Oracle Spatial sample programs: <ul style="list-style-type: none"> ■ X11 for Mac OS X and X11 SDK for Mac OS X ■ OpenMotif 2.2.3 or higher <p>Note: This software is required only to build the sample programs. For information about obtaining OpenMotif, see the following Web site: http://darwinports.opendarwin.org/</p>

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

1. To determine the version of the operating system, enter the following command:

```
# sw_vers
```

Note: Only Mac OS X Server is supported.

2. To determine whether the required version of `gcc` is installed, enter the following command:

```
# gcc -v
```

This command returns output similar to the following:

```
Reading specs from /usr/libexec/gcc/darwin/ppc/3.3/specs
Thread model: posix
gcc version 3.3 20030304 (Apple Computer, Inc. build 1635) (-mlongcall)
```

If the required version (including the date and build number) or a later version is not installed, download and install either of the following:

- The Dec 2003 gccLongBranch Tools update
- Xcode Tools 1.2 or higher, which includes the correct version of gcc

You can download this software from the Apple Developer Connection Web site:

<http://connect.apple.com>

After you log in, choose **Download Software**, then **Developer Tools** to display the download page.

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, lightweight directory access protocol (LDAP) or 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=/Volumes/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, use one of the following methods to create the Oracle Inventory group:

Note: Oracle recommends that you use the group name `oinstall` for the Oracle Inventory group.

- Use Workgroup Manager (`/Applications/Server/Workgroup Manager`) to create a group called `oinstall`.
- Use command-line utilities to create the Oracle Inventory group:
 1. Enter the following command to list existing groups and group IDs (GIDs):

```
# nireport . /groups gid name | more
```

2. Identify an unused GID for the new Oracle Inventory group.

3. Enter commands similar to the following to create the `oinstall` group:

```
# nicl . -create /groups/oinstall
# nicl . -append /groups/oinstall gid gid_number
# nicl . -append /groups/oinstall passwd "*" 
```

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 OSDBA 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. In the following procedures, use the group name `dba` unless a group with that name already exists.

- Use Workgroup Manager to create a group called `dba`.
- Use command-line utilities to create the OSDBA group:
 1. Enter the following command to list existing groups and GIDs:

```
# nireport . /groups gid name | more
```

2. Identify an unused GID for the new OSDBA group.

3. Enter commands similar to the following to create the `dba` group:

```
# nicl . -create /groups/dba
# nicl . -append /groups/dba gid gid_number
```



```
# nicl . -append /groups/dba passwd "*"

```

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. In the following procedures, use the group name `oper` unless a group with that name already exists.

- Use Workgroup Manager to create a group called `oper`.
- Use command-line utilities to create the OSOPER group:
 1. Enter the following command to list existing groups and GIDs:

```
# nireport . /groups gid name | more

```

2. Identify an unused GID for the new OSOPER group.
3. Enter commands similar to the following to create the `oper` group:

```
# nicl . -create /groups/oper
# nicl . -append /groups/oper gid gid_number
# nicl . -append /groups/oper passwd "*"

```

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 (`gid`) 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 6-11.
- To modify an existing user, see the "[Modify an Existing Oracle Software Owner User](#)" section on page 6-11.
- 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. In the following procedures, use the user name `oracle` unless a user with that name already exists.

- Use Workgroup Manager to create a user with the following properties:

Property	Recommended Value
Name	Oracle Software Owner
Short name	<code>oracle</code>
Password	Specify and verify a password for the <code>oracle</code> user
Primary group	<code>oinstall</code>
Other groups	The OSDBA group (<code>dba</code>), and optionally the OSOPER group (<code>oper</code>) if you created it
Default shell	Any supported shell, for example <code>/bin/tcsh</code> or <code>/bin/bash</code>
Home directory	A home directory consistent with other user home directories, for example <code>/Users/oracle</code>

- Use command-line utilities to create the Oracle software owner user:
 1. Enter the following command to list existing users and user IDs (UIDs):


```
# nireport . /users uid name | more
```
 2. Identify an unused UID for the new Oracle software owner user.
 3. Enter commands similar to the following to create the `oracle` user:


```
# nicl . -create /users/oracle
# nicl . -append /users/oracle uid uid_number
# nicl . -append /users/oracle gid oinstall_gid
# nicl . -append /users/oracle shell /bin/bash
# nicl . -append /users/oracle home /Users/oracle
# nicl . -append /users/oracle realname "Oracle software owner"
```
 4. Enter commands similar to the following to add the `oracle` user to the OSDBA and OSOPER groups (`dba` and `oper`):


```
# nicl . -append /groups/dba users oracle
# nicl . -append /groups/oper users oracle
```
 5. Create the home directory for the `oracle` user and change the owner and group on that directory:

```
# mkdir /Users/oracle
# chown oracle:oinstall /Users/oracle
```

6. Set the password for the `oracle` user:

```
# passwd oracle
```

See the ["Verifying That the UNIX User nobody Exists"](#) section on page 6-11 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 use one of the following methods to modify it:

- Use Workgroup Manager to modify the `oracle` user, specifying `oinstall` as the Primary group and the appropriate OSDBA or OSOPER groups, for example `dba` and `oper`, as secondary groups.

- Use command-line utilities to modify the `oracle` user:

1. Enter the following command to determine the GID of the `oinstall` group:

```
# nireport . /groups/oinstall gid name
```

2. Enter the following command to change the GID of the `oracle` user to the GID of the `oinstall` group:

```
# nicl . -delete /users/oracle gid
# nicl . -append /users/oracle gid oinstall_gid
```

3. If necessary, enter the following commands to add the `oracle` user to the OSDBA and OSOPER groups:

```
# nicl . -append /groups/dba users oracle
# nicl . -append /groups/oper users oracle
```

Verifying That the UNIX User nobody Exists

Note: The `nobody` user exists by default on all Mac OS X systems.

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, create it as follows:

a. Enter the following command to determine whether the `nobody` group exists:

```
# nireport . /groups name gid | grep nobody
```

b. If the `nobody` group does not exist, enter commands similar to the following to create it:

```
# nicl . -create /groups/nobody
# nicl . -append /groups/nobody gid -2
```

```
# nicl . -append /groups/nobody passwd "*"
```

- c. Enter commands similar to the following to create the nobody user:

```
# nicl . -create /users/nobody
# nicl . -append /users/nobody uid -2
# nicl . -append /users/nobody gid -2
# nicl . -append /users/nobody shell /usr/bin/false
# nicl . -append /users/nobody home /var/empty
# nicl . -append /users/nobody realname "Unprivileged User"
# nicl . -append /users/nobody passwd "*"

```

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 or if all cluster nodes use the same parent NetInfo database, 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 identical users and groups on the other nodes, use one of the following methods:

- Use Workgroup Manager to create or modify existing groups and users on the other cluster nodes:
 1. Determine whether the `oinstall`, `dba`, and `oper` groups exist and identify their group IDs (GIDs):

Note: If the GID used by a group on the first node is unavailable on another node, you must change the GID on all of the nodes, including the first node, to a GID that is available on all nodes.

- If the GID for an existing `oinstall`, `dba`, or `oper` group is not the same as its GID on the other nodes, modify the GID so that it is the same on all nodes.
 - If an `oinstall`, `dba`, or `oper` group does not exist on a node, create it and specify the same GID that it has on the other nodes.
2. Determine whether the `oracle` user exists and identify its primary and secondary (Other) groups:

Note: If the UID used by the `oracle` user on the first node is unavailable on another node, you must change the UID on all of the nodes, including the first node, to a UID that is available on all nodes.

- If the `oracle` user exists but its UID is not the same as on the other nodes, its primary group is not `oinstall`, or it is not a member of the `dba` group (and optionally the `oper` group), modify the user properties as required.
- If the `oracle` user does not exist, create it by following the instructions in the "[Create a New Oracle Software Owner User](#)" section on page 6-10.

Make sure that the UID that you specify for the user is the same as its UID on the other nodes.

- Use command-line utilities to create or modify existing groups and users on the other cluster nodes:

1. Enter the following command to list existing groups and GIDs:

```
# nireport . /groups gid name | more
```

From the output, determine whether the `oinstall`, `dba`, and `oper` groups exist:

Note: If the GID used by a group on the first node is unavailable on another node, you must change the GID on all of the nodes, including the first node, to a GID that is available on all nodes.

- If a group exists, verify that it has the same GID as the corresponding group on the first node. If it uses a different GID, enter commands similar to the following to change the GID:

```
# nicl . -delete /groups/group_name gid
# nicl . -append /groups/group_name gid new_gid
```

- If a group does not exist, verify that the GID used by the group on the first node is unused, then enter commands similar to the following to create it:

```
# nicl . -create /groups/group_name
# nicl . -append /groups/group_name gid gid_number
# nicl . -append /groups/group_name passwd "*"
```

2. Enter the following commands to ensure that the `oracle` user is a member of the correct groups:

```
# nicl . -merge /groups/dba users oracle
# nicl . -merge /groups/oper users oracle
```

3. Enter the following command to identify whether the `oracle` user exists and to display the UIDs of existing users:

```
# nireport . /users uid gid name | more
```

From the output, determine whether the `oracle` user exists:

Note: If the UID used by the `oracle` user on the first node is unavailable on another node, you must change the UID on all of the nodes, including the first node, to a UID that is available on all nodes.

- If the `oracle` user exists, but its UID is different from the UID on the first node, enter commands similar to the following to change the UID:

```
# nicl . -delete /users/oracle uid
# nicl . -append /users/oracle uid new_uid
```

- If the `oracle` user exists, but its GID is different from the GID for the `oinstall` group, enter commands similar to the following to change the GID:

```
# nicl . -delete /users/oracle gid
# nicl . -append /users/oracle gid oinstall_gid
```

- If the `oracle` user does not exist, verify that the UID used by the `oracle` user on the first node is unused, then enter commands similar to the following to create it:

```
# nicl . -create /users/oracle
# nicl . -append /users/oracle uid uid_number
# nicl . -append /users/oracle gid oinstall_gid
# nicl . -append /users/oracle shell /bin/bash
# nicl . -append /users/oracle home /Users/oracle
# nicl . -append /users/oracle realname "Oracle software owner"
```

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.

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 run the Installer, 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.

If you see any other messages or text, apart from the date, the installation might fail. Make any changes required to ensure that only the date is displayed when you enter these commands.

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

Note: 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.

If you are installing the RAC option, you must set the kernel parameters and shell limits on all cluster nodes.

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	Value
kern.corefile	core
kern.sysv.semmni	87381
kern.sysv.semmns	87381
kern.sysv.semmnu	87381
kern.sysv.semmsl	87381
kern.sysv.semume	10
kern.sysv.maxproc	2068
kern.sysv.maxprocperuid	2068

Parameter	Value
net.inet.udp.maxdgram	65536
net.inet.udp.recvspace	65536

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 table 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
kern.corefile	# /usr/sbin/sysctl -a grep corefile
kern.sysv.semni, kern.sysv.semns, kern.sysv.semnu, kern.sysv.semmsl, kern.sysv.semume	# /usr/sbin/sysctl -a grep sem
kern.maxproc, kern.maxprocperuid	# /usr/sbin/sysctl -a grep maxproc
net.inet.udp.maxdgram, net.inet.udp.recvspace	# /usr/sbin/sysctl -a grep udp

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.

```
kern.corefile=core
kern.sysv.semni=87381
kern.sysv.semns=87381
kern.sysv.semnu=87381
kern.sysv.semmsl=87381
kern.sysv.semume=10
kern.maxproc=2068
kern.maxprocperuid=2068
net.inet.udp.maxdgram=65536
net.inet.udp.recvspace=65536
```

By specifying the values in the `/etc/sysctl.conf` file, they persist when you restart the system.

- b. For each parameter that you must modify, enter a command similar to the following to change the current value:

```
# /usr/sbin/sysctl -w parameter_name=value
```

For example:

```
# /usr/sbin/sysctl -w kern.maxproc=2068
```

Set Shell Limits

To improve the performance of the software on Mac OS X systems, you must increase the following shell limits:

Shell Limit	Hard Limit
Maximum number of open file descriptors	65536
Maximum number of processes available to a single user	2068

To increase the shell limits:

1. Change directory to the following directory:
2. Open the `IPServices` script using any text editor.
3. Add the following `ulimit` commands to the start of the `StartService` function in this script:

```
ulimit -Hu 2068
ulimit -Su 2068
ulimit -Hn 65536
ulimit -Sn 65536
```

4. Add the same `ulimit` commands to the `/etc/rc` script, before the `SystemStarter` command.

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 `/Volumes/u01` for the mount point directory. However, you could choose another mount point directory, `/Volumes/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:

```
/Volumes/u01/app/oracle
/Volumes/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. 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. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
/Volumes/u01/crs/oracle/product/10.1.0/crs
```

Note: Because you must change the permissions of all of the parent directories of the CRS home directory after installing the software to grant write access only to the `root` user, the CRS home directory must not be a subdirectory of the Oracle base 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=/Volumes/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, `/Volumes/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:

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

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

Note: If possible, choose a directory path similar to the first (`/Volumes/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 the same path on the other nodes.
- It must have at least 3 GB of free disk space.

To determine the free disk space on the file system where the Oracle base directory is located, enter the following command:

```
# df -h oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, see the "[Create the CRS Home Directory](#)" section on page 6-22.

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 3 GB of free disk space.

To identify an appropriate file system, follow these steps:

1. Use the `df -h` 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.

For RAC installations, 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/oracle_sw_owner
# chmod -R 775 /mount_point/app/oracle_sw_owner
```

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

```
/Volumes/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.

Create the CRS Home Directory

You must create a CRS home directory before installing Oracle CRS. You can create the CRS home directory on the same file system as the Oracle base directory, or you can choose a different file system for this directory. If you choose the same file system as the Oracle base directory, you must not use a subdirectory of the Oracle base directory as the CRS home directory.

To identify an appropriate file system, follow these steps:

1. Use the `df -h` command to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has at least 1 MB of free disk space.

If you are using the same file system for the Oracle base directory, this 1 MB of disk space is additional to the free disk space requirement that you identified previously.

Note: The file system must be a local file system.

The path to the CRS home 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 CRS home 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/crs/oracle_sw_owner/product/10.1.0/crs
# chown -R oracle:oinstall /mount_point/crs/oracle_sw_owner
# chmod -R 775 /mount_point/crs/oracle_sw_owner
```

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

```
/Volumes/u01/crs/oracle/product/10.1.0/crs
```

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 installing Oracle CRS, set the `ORACLE_HOME` environment variable to specify this directory.

Note: During the installation, before you run the `root.sh` script, you must change the permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories.

Choose a Storage Option for Oracle CRS, Database, and Recovery Files

The following table shows the storage options supported for storing Oracle Cluster Ready Services (CRS) files, Oracle database files, and Oracle database recovery files. Oracle database files include datafiles, control files, redo log files, the server parameter file, and the password file. Oracle CRS files include the Oracle Cluster Registry (OCR) and the CRS voting disk.

For all installations, you must choose the storage option that you want to use for Oracle CRS files and Oracle database files. If you want to enable automated backups during the installation, you must also choose the storage option that you want to use for recovery files (the flash recovery area). You do not have to use the same storage option for each file type.

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>

Storage Option	File Types Supported		
	CRS	Database	Recovery
Automatic Storage Management	No	Yes	Yes
Shared raw disk partitions	Yes	Yes	No

Use the following guidelines when choosing the storage option that you want to use for each file type:

- You can choose any combination of the supported storage options for each file type.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For Standard Edition installations, ASM is the only supported storage option for database or recovery files.
- You cannot use Automatic Storage Management to store Oracle CRS files, because these files must be accessible before any Oracle instance starts.

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 6-24.
- To use raw devices for Oracle CRS or database file storage, see the "[Configure Raw Partitions](#)" section on page 6-29.

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

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, recovery files, or both.

Note: For RAC installations, if you choose to enable automated backups, 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 on 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 database files 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 method 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 database files 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	Database Files	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

For RAC installations, 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 ASM instance is already running 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 protect against 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 6-28.

Using an Existing ASM Disk Group

If you want to store either database or recovery files in an existing ASM disk group, 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 method 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 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

Note: On Mac OS X systems, Oracle recommends that you create a single whole-disk partition on each disk that you want to use.

To configure disks for ASM using raw devices, follow these steps:

1. If necessary, install the shared disks that you intend to use for the disk group and restart the system.
2. Enter the following command to display information about the disks attached to the system:

```
# diskutil list
```

From the display, identify the name of the disk devices that you want to use for the ASM disk group.

3. Enter the following command to start the Apple partition table editor:

```
# /usr/sbin/pdisk
```

4. To create a single whole-disk partition on each device that you want to include in the disk group, repeat the following steps:

- a. Enter `e` to edit the partition table.
- b. Enter the device name of the disk that you want to partition, for example `/dev/rdisk4`.
- c. Enter `p` to display the partition table for that disk.
- d. If the disk does not have a partition table, enter `i` to create one.

Press Enter at each prompt to accept the default values for block and device sizes, then enter `p` again to display the new partition table. The output should be similar to the following:

```
/dev/rdisk4 map block size=512
#:                type name                length  base  ( size )
 1:  Apple_partition_map Apple                63 @ 1
 2:                Apple_Free Extra          20479936 @ 64  ( 100.0G)
```

```
Device block size=512, Number of Blocks=204800000
DeviceType=0x0, DeviceId=0x0
```

- e. Enter `C` to create a new disk partition and specify the partition type.
- f. Enter the following values for the new partition:

Prompt	Recommended Response
First block	Enter the block number associated with the start (base) of the free space on the disk, for example 64.

Prompt	Recommended Response
Length in blocks	Enter the length in blocks (the <code>length</code> value) of the free space on the disk, for example 20479936.
Name of partition	Enter a name for the partition, for example <code>ASM_Disk1</code> .
Type of partition	Enter <code>Apple_Label</code> .

- g. Enter `p` again, to display the new partition table.
The type displayed for the new partition should be `Apple_Label`.
 - h. Enter `w`, then `y` to write the new partition table to the disk.
 - i. Enter `q` to return to the top-level menu.
5. When you have finished creating partitions, enter `q` to quit from `pdisk`.
 6. For each disk device that you want to use, enter commands similar to the following to create new device files for the disk partition that you created, with the correct owner, group, and permissions:

```
# /sbin/disklabel -create /dev/diskns2 owner-uid=oracle owner-gid=dba \
owner-mode=0660
```

In this example, the device name specifies slice 2, which represents the whole-disk partition that you created on `diskn`. This command creates additional block and character device files for the partition that you specified, owned by the `oracle` user and `dba` group, and with permissions 660:

```
/dev/diskns2s1
/dev/rdiskns2s1
```

When you specify disk devices for ASM, you must specify the character device file name (`/dev/rdiskns2s1`).

7. On the other cluster nodes, enter commands similar to the following to create the new device files for each disk device on that node:

```
# /usr/sbin/pdisk /dev/diskn -dump
```

Note: If the nodes are configured differently, the disk device names might be different on some nodes. Make sure that you specify the correct device names in this command on each node.

Configure Raw Partitions

Note: If you are using ASM for database file storage, you need only create raw partitions for the Oracle CRS files.

Table 6–1 lists the number, size, and sample partition names for the partitions that you must create for database files. Table 6–2 lists the number, size, and sample partition names for the partitions that you must create for CRS files.

Table 6–1 Raw Partitions Required for Database Files on Mac OS X

Number	Size (MB)	Purpose and Sample Partition Device Name
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 <i>n</i> tablespace (One tablespace for each instance) <i>dbname_undotbsn_raw_500m</i>
1	250	TEMP tablespace <i>dbname_temp_raw_250m</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 <i>dbname_redon_m_raw_120m</i>
2	110	First and second control files <i>dbname_control{1 2}_raw_110m</i>
1	5	Server parameter file (SPFILE) <i>dbname_spfile_raw_5m</i>
1	5	Password file <i>dbname_pwdfile_raw_5m</i>

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.

Table 6–2 Raw Partitions Required for CRS Files on Mac OS X

Number	Size (MB)	Purpose and Sample Partition Name
1	100	Oracle Cluster Registry <i>ora_ocr_raw_100m</i> 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.
1	20	Oracle CRS voting disk <i>ora_vote_raw_20m</i> 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.

The following subsections describe how to configure raw partitions on Mac OS X.

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 raw partitions and restart the system.
2. Enter the following command to display information about the disks attached to the system:

```
# diskutil list
```

From the display, identify the name of the disk devices that you want to use for the raw partitions.

3. Enter the following command to start the Apple partition table editor:

```
# /usr/sbin/pdisk
```

4. To create partitions on each device that you want to use, repeat the following steps:

- a. Enter `e` to edit the partition table.
- b. Enter the device name of the disk that you want to partition, for example `/dev/rdisk4`.
- c. Enter `p` to display the partition table for that disk.
- d. If the disk does not have a partition table, enter `i` to create one.

Press Enter at each prompt to accept the default values for block and device sizes, then enter `p` again to display the new partition table. The output should be similar to the following:

```
/dev/rdisk4 map block size=512
#:          type name                length  base   ( size )
1:  Apple_partition_map Apple          512 @ 1
2:          Apple_Free  Extra        209750425 @ 513   ( 100.0G)
```

```
Device block size=512, Number of Blocks=204800000
DeviceType=0x0, DeviceId=0x0
```

- e. Enter `C` to create new disk partitions and specify their type.

Enter the following values for each new partition that you want to create:

Note: Do not create more than 15 partitions on any disk. See the `pdisk` man page for more information.

Prompt	Recommended Response
First block	Enter the block number associated with the start (base) of the free space on the disk, for example 513.
Length in blocks	Enter the size of the partition that you want to create. For example, enter 500m to create a 500 MB partition.
Name of partition	Enter a name for the partition, for example <code>dbname_system_raw_500m</code> . In this example, <code>dbname</code> is the name that you want to call the database, <code>system</code> is the database object identifier, and 500m is the raw partition size.

Prompt	Recommended Response
Type of partition	Enter <code>Apple_Label</code> .

- f. Enter `p` again, to display the new partition table.
The type displayed for the new partitions should be `Apple_Label`.
 - g. Enter `w`, then `y` to write the new partition table to the disk.
 - h. Enter `q` to return to the top-level menu.
5. When you have finished creating partitions, enter `q` to quit from `pdisk`.
 6. For each partition that you created, enter commands similar to the following to create new device files for that partition, with the correct owner, group, and permissions:

- OCR device:

```
# /sbin/disklabel -create /dev/disknsx owner-uid=root \
owner-gid=oinstall owner-mode=0640 dev-name=ora_ocr_raw_100m
```

- Other devices:

```
# /sbin/disklabel -create /dev/disknsx owner-uid=oracle \
owner-gid=dba owner-mode=0660 dev-name=sample_partition_name
```

For example:

```
# /sbin/disklabel -create /dev/disk4s3 owner-uid=oracle \
owner-gid=dba owner-mode=0660 dev-name=orcl_system_raw_500m
```

In these examples, the device name specifies partition (slice) *x* on disk *n*. The disk label command creates additional block and character device files for the partition that you specified, with the specified user, group, and permissions. For example

```
/dev/orcl_system_raw_500m
/dev/orcl_system_raw_500m
```

When you specify raw devices in the raw device mapping file or during the installation, you must specify the character device file name (for example `/dev/orcl_system_raw_500m`).

7. On the other cluster nodes, enter commands similar to the following for each disk to create the new device files for the new partitions on that disk:

```
# /usr/sbin/pdisk /dev/diskn -dump
```

Note: If the nodes are configured differently, the disk device names might be different on some nodes. Make sure that you specify the correct device names in this command on each node.

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

- C shell:

```
% setenv ORACLE_BASE /Volumes/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/rdbname_system_raw_500m
sysaux=/dev/rdbname_sysaux_raw_800m
example=/dev/rdbname_example_raw_160m
users=/dev/rdbname_users_raw_120m
temp=/dev/rdbname_temp_raw_250m
undotbs1=/dev/rdbname_undotbs1_raw_500m
undotbs2=/dev/rdbname_undotbs2_raw_500m
redo1_1=/dev/rdbname_redo1_1_raw_120m
redo1_2=/dev/rdbname_redo1_2_raw_120m
redo2_1=/dev/rdbname_redo2_1_raw_120m
redo2_2=/dev/rdbname_redo2_2_raw_120m
control1=/dev/rdbname_control1_raw_110m
control2=/dev/rdbname_control2_raw_110m
spfile=/dev/rdbname_spfile_raw_5m
pwdfile=/dev/rdbname_pwdfile_raw_5m
```

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 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.

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 -auxwww | 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.

Note: If no Oracle Net listeners are running, see the "[Configure the oracle User's Environment](#)" section on page 6-35 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 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 ORACLE_BASE environment variable.

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

1. Start a new terminal session.
2. Use the `ssh`, `rlogin`, or `telnet` command to log in to the system where you want to install the software as the `oracle` user.

Note: You must complete this step even if you are installing the software on the local system.

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

```
$ echo $SHELL
```

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

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

```
$ vi .profile
```

Note: If the oracle user's shell is the Bash shell, open the `.bash_profile` file, if it exists.

- C shell (csh or tcsh):


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


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


```
$ . ~/.bash_profile
```
 - Bourne, Bash, or Korn shell:


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


```
% source ~/.login
```
- 9. 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 -h` command to identify a suitable file system with sufficient free space.
 - b. If necessary, as a user that is a member of the `admin` group, enter commands similar to the following to create a temporary directory on the file system that you identified, and set the appropriate permissions on the directory:


```
$ sudo mkdir /mount_point/tmp
          $ sudo chmod a+wr /mount_point/tmp
```
 - c. As the `oracle` user, 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
```
- 10. Enter commands similar to the following to set the `ORACLE_BASE` environment variable:

- Bourne, Bash, or Korn shell:

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

- C shell:

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

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

11. If you are using raw devices for database file 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
```

12. Enter the following command to ensure that the `ORACLE_HOME` and `TNS_ADMIN` environment variables are not set:

- Bourne, Bash, or Korn shell:

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

- C shell:

```
% unsetenv ORACLE_HOME
% unsetenv TNS_ADMIN
```

Note: If the `ORACLE_HOME` environment variable is set, the Installer uses the value that 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.

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

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

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

Pre-Installation Tasks for RAC on Solaris

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer (OUI). 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
- Create the CRS Home Directory
- Choose a Storage Option for Oracle CRS, Database, and Recovery Files
- Create Directories for Oracle CRS, Database, or Recovery Files
- Configure Disks for Automatic Storage Management
- Configure Raw Partitions or 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-mode 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 silent-mode installations, refer to [Appendix B](#).

- 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 are not installing the software on the local system, use the `ssh`, `rlogin`, or `telnet` commands to connect to the system where you want to install the software:

```
$ 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, refer to 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

Each 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 requirement 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 each 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, refer to 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 whether the system architecture can run the software, enter the following command:

```
# /bin/isainfo -kv
```

Verify that this command displays the correct output:

- SPARC systems:
64-bit `sparcv9` kernel modules
- x86 systems:
32-bit `i386` 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, refer to 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)
 - (SPARC systems only) Remote shared memory (RSM) using peripheral component interconnect-scalable coherent interface (PCI-SCI) adapters and SCI switches

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

Note: If you are using Sun Cluster 3.1 (SPARC systems only), you must use the private host name and IP address configured for each node during the Sun Cluster installation. The default private host name for each Sun Cluster node is `clusternodeid-priv`. Do **not** specify this host name or IP address in the `/etc/hosts` or `/etc/inet/hosts` file.

If you are not using Sun Cluster 3.1, you must configure private host names and IP addresses for each node.

Oracle recommends that you use non-routable IP addresses for the private interfaces, for example: `10.*.*` or `192.168.*.*`. If you are not using Sun Cluster 3.1, 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	143.47.43.104	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.

Note: If you are using Sun Cluster 3.1, go to step 6

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    racl-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.

6. If you are using Sun Cluster 3.1, enter the following command to display private host names:

```
# /usr/cluster/bin/scconf -p
```

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

Check the Software Requirements

Refer to 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.

Checking the Software Requirements

Check that the required software and patches are installed on the system.

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>SPARC systems:</p> <p>Solaris 8 or Solaris 9, 64-bit</p> <p>x86 systems:</p> <p>Solaris 9</p> <hr/> <p>Operating system packages:</p> <p>SPARC and x86 systems:</p> <p>SUNWarc SUNWbtool SUNWhea SUNWlibC SUNWlibm SUNWlibms SUNWsprout SUNWtoo SUNWilof SUNWilcs SUNWil5cs SUNWxwfont</p> <p>SPARC systems only:</p> <p>SUNWsprout</p> <p>Note: You might also require additional font packages for Java, depending on your locale. Refer to the following Web site for more information: 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	<p>Sun Cluster 3.1 (SPARC systems only)</p> <p>Note: Sun Cluster is supported for use with RAC on SPARC systems but it is not required.</p>
Oracle Messaging Gateway	<p>Oracle Messaging Gateway supports the integration of Oracle Streams Advanced Queuing (AQ) with the following software:</p> <ul style="list-style-type: none"> ■ IBM WebSphere MQ V5.3, client and server (SPARC): mqm ■ IBM MQSeries V5.1, client and server (x86) ■ TIBCO Rendezvous 7.2 (SPARC only)
PL/SQL native compilation, 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)

Installation Type or Product	Requirement
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 (SPARC) or 1.3.1_11 (x86) 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 (SPARC) or 1.4.2_04 (x86) is installed with this release.</p>
Oracle Transparent Gateway for Sybase (SPARC only)	<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>
Oracle Transparent Gateway for Informix (SPARC only)	<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 Ingres (SPARC only)	Ingres II version 2.0, 2.5, or 2.6
Oracle Transparent Gateway for Teradata (SPARC only)	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
5.9
```

In this example, the version shown is Solaris 9 (5.9). If necessary, refer to your operating system documentation for information about upgrading the operating system.

2. To determine whether the required packages are installed, enter a command similar to the following:

```
# pkginfo -i SUNWarc SUNWbtool SUNWhea SUNWlibC SUNWlibm SUNWlibms SUNWsprout \
SUNWsprprox SUNWttoo SUNWwilof SUNWilcs SUNWi15cs SUNWxwfn
```

If a package is not installed, then install it. Refer to 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 (SPARC only):</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: IIIM 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 & /.../exacctsypatch ■ 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/libdhcpagent.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 (SPARC):</p> <ul style="list-style-type: none"> ■ 112233-11, SunOS 5.9: Kernel Patch ■ 111722-04, SunOS 5.9: Math Library (libm) patch <p>Patches for Solaris 9 (x86):</p> <ul style="list-style-type: none"> ■ 111713-06, SunOS 5.9_x86: Shared library patch for C++ ■ 111728-03, SunOS 5.9_x86: Math Library (libm) patch ■ 112234-12, SunOS 5.9_x86: Kernel Patch ■ 113986-08, SunOS 5.9_x86: linker Patch ■ 115114-02, SunOS 5.9_x86: Patch for assembler ■ 116013-02, SunOS 5.9_x86: ps utility patch (Available only from your local Sun solution center)

Installation Type or Product	Requirement
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	<p>Patches for Solaris 8 and Solaris 9 (SPARC):</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 <p>Patches for Solaris 9 (x86):</p> <ul style="list-style-type: none"> ■ 112756-10, Patch for S1S8CC Sun Compiler Common 7.1 compiler backend_x86 ■ 112759-03, dbx 7.1_x86: Patch for S1S8CC Debugger ■ 112761-07, C 5.5_x86: Patch for S1S8CC C compiler ■ 113819-07, C++ 5.5_x86: Patch for S1S8CC C++ compiler ■ 113821-01, Compiler Common 7.1_x86: Patch for S1S8CC Math Libraries ■ 113825-02, ILD 4.1_x86: Patch for S1S8CC Incremental Linker
Real Application Clusters	<p>Sun Cluster patches for Solaris 8 (SPARC):</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 (SPARC):</p> <p>113801: Sun Cluster 3.1: Core/Sys Admin Patch</p>
Oracle Messaging Gateway	<p>Corrective service diskettes (CSDs) for WebSphere MQ:</p> <p>SPARC systems:</p> <p>CSD06 or later for WebSphere MQ V5.3</p> <p>x86 systems:</p> <ul style="list-style-type: none"> ■ CSD09 or later for MQSeries Server V5.1 ■ MQSeries Client for Sun Solaris, Intel Platform Edition- V5.1 SupportPac MACE

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 or a SupportPac for WebSphere MQ or MQSeries, refer to 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 refer to 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, enter the following command to create the `oinstall` group:

```
# /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 OSDBA 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, enter the following command to create the dba group. Use 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, enter the following command to create the oper group. Use 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. Refer to 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, refer to the [Verifying that the UNIX User nobody Exists](#) section on page 7-14.
- To modify an existing user, refer to the "[Modify an Existing Oracle Software Owner User](#)" section on page 7-14.
- To create a new user, refer to 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. Use 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
```

Refer to the "[Verifying that the UNIX User nobody Exists](#)" section on page 7-14 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, enter a command similar to the following to modify it. Specify 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 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
```

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`. Refer to 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 node1-public date
$ rsh node2-public 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 and Shell Limits

Refer to the following section for information about configuring the system kernel parameters and shell limits.

Configuring Kernel Parameters

Note: The kernel parameter 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. Refer to your operating system documentation for more information about tuning kernel parameters.

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
noexec_user_stack (SPARC only)	1
semsys:seminfo_semmni	100
semsys:seminfo_semmns	1024
semsys:seminfo_semmsl	256
semsys:seminfo_sevmx	32767
shmsys:shminfo_shmmax	4294967295
shmsys:shminfo_shmmin ¹	1
shmsys:shminfo_shmmni	100
shmsys:shminfo_shmse 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_stack /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_semmns1=256
set semsys:seminfo_semvms=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 restart 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
```


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, a supported cluster file system, or on an NFS file system on a certified NAS device.

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 or an NFS file system on a NAS device, you must specify a path for the Oracle Inventory directory on a local file system. The Oracle base directory must be 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. Oracle recommends that you specify a path similar to the following for the CRS home directory:

```
/u01/crs/oracle/product/10.1.0/crs
```

Note: Because you must change the permissions of all of the parent directories of the CRS home directory after installing the software to grant write access only to the `root` user, the CRS home directory must not be a subdirectory of the Oracle base 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 /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 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 3 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 4 GB

To determine the free disk space on the file system where the Oracle base directory is located, enter the following command:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, refer to the ["Create the CRS Home Directory"](#) section on page 7-23.

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, refer to 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 3 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 4 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, 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/oracle_sw_owner
# chmod -R 775 /mount_point/app/oracle_sw_owner
```

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.

Create the CRS Home Directory

You must create a CRS home directory before installing Oracle CRS. You can create the CRS home directory on the same file system as the Oracle base directory, or you can choose a different file system for this directory. If you choose the same file system as the Oracle base directory, you must not use a subdirectory of the Oracle base directory as the CRS home directory.

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 at least 1 MB of free disk space.

If you are using the same file system for the Oracle base directory, this 1 MB of disk space is additional to the free disk space requirement that you identified previously.

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 CRS home 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 CRS home 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/crs/oracle_sw_owner/product/10.1.0/crs
# chown -R oracle:oinstall /mount_point/crs/oracle_sw_owner
# chmod -R 775 /mount_point/crs/oracle_sw_owner
```

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

```
/u01/crs/oracle/product/10.1.0/crs
```

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 installing Oracle CRS, set the `ORACLE_HOME` environment variable to specify this directory.

Note: During the installation, before you run the `root.sh` script, you must change the permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories.

Choose a Storage Option for Oracle CRS, Database, and Recovery Files

The following table shows the storage options supported for storing Oracle Cluster Ready Services (CRS) files, Oracle database files, and Oracle database recovery files. Oracle database files include datafiles, control files, redo log files, the server parameter

file, and the password file. Oracle CRS files include the Oracle Cluster Registry (OCR) and the CRS voting disk.

For all installations, you must choose the storage option that you want to use for Oracle CRS files and Oracle database files. If you want to enable automated backups during the installation, you must also choose the storage option that you want to use for recovery files (the flash recovery area). You do not have to use the same storage option for each file type.

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

<http://metalink.oracle.com>

Storage Option	File Types Supported		
	CRS	Database	Recovery
Automatic Storage Management	No	Yes	Yes
Cluster file system	Yes	Yes	Yes
Note: Requires a supported cluster file system			
NFS file system	Yes	Yes	Yes
Note: Currently supported only with Fujitsu PRIMECLUSTER and a certified NAS device (SPARC only)			
Shared raw logical volumes (SPARC only)	Yes	Yes	No
Shared raw partitions	Yes	Yes	No

Use the following guidelines when choosing the storage options that you want to use for each file type:

- You can choose any combination of the supported storage options for each file type as long as you satisfy any requirements listed for the chosen storage options.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For Standard Edition installations, ASM is the only supported storage option for database or recovery files.
- You cannot use Automatic Storage Management to store Oracle CRS files, because these files must be accessible before any Oracle instance starts.

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

- To use a file system for Oracle CRS, database, or recovery file storage, refer to the "[Create Directories for Oracle CRS, Database, or Recovery Files](#)" section on page 7-25.
- To use ASM for database or recovery file storage, refer to the "[Configure Disks for Automatic Storage Management](#)" section on page 7-28.
- To use raw devices (partitions or logical volumes) for Oracle CRS or database file storage, refer to the "[Configure Raw Partitions or Raw Logical Volumes](#)" section on page 7-34.

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 on shared disks, for example:
 - A supported cluster file system
 - 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 that 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 on shared disks, for example:
 - A supported cluster file system
 - 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 that 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 that 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 on shared disks, for example:
 - A supported cluster file system
 - 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, refer to the *Oracle Backup and Recovery Basics* manual.

- The default path suggested by the Installer for the flash recovery area 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 that 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: The mount point that you use for the file system must be identical on each node. Make sure that the file systems are configured to mount automatically when a node restarts.

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
Database files	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, refer to one of the following sections:
 - [Configure Disks for Automatic Storage Management](#) on page 7-28
 - [Configure Raw Partitions or Raw Logical Volumes](#) on page 7-34Otherwise, refer to the "[Verify that the Cluster Software is Configured and Running](#)" section on page 7-44.

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](#)
- [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. Refer to the *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, recovery files, or both.

Note: You do not have to use the same storage mechanism for database files and recovery files. You can use the file system for one file type and ASM for the other.

For RAC installations, if you choose to enable automated backups and you do not have a shared 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 on 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 database files 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 method 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 database files 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	Database Files	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

For RAC installations, 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 ASM instance is already running 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 protect against 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, refer to the ["Configuring Disks for ASM"](#) section on page 7-32.

Using an Existing ASM Disk Group

If you want to store either database or recovery files in an existing ASM disk group, 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 method 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 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 on Solaris, follow these steps:

1. If necessary, install the shared disks that you intend to use for the disk group and restart the system.
2. On SPARC systems only, enter the following commands to identify devices that are part of a logical volume manager (LVM) disk group:

Note: The following command displays information about VERITAS Volume Manager disks. If you use a different LVM, refer to 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 an LVM disk group. Do not use disks that are part of an LVM disk group.

3. To create or identify the disk slices (partitions) that you want to include in the ASM disk group, follow these steps:
 - a. To list the disks attached to the system, 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`).

- b. Enter the number corresponding to the disk that you want to use.
- c. On x86 systems only, use the `fdisk` command to create a Solaris partition on the disk if one does not already exist.

Solaris `fdisk` partitions must start at cylinder 1, not cylinder 0. If you create a new `fdisk` partition, you must label the disk before continuing.

- d. Enter the `partition` command, followed by the `print` command to display the partition table for the disk that you want to use.
- e. If necessary, create a single whole-disk slice, starting at cylinder 1.

---

**Note:** To prevent ASM from overwriting the partition table, you cannot use slices that start at cylinder 0 (for example, slice 2).

---

- f. Make a note of the number of the slice that you want to use.
  - g. If you modified a partition table or created a new one, enter the `label` command to write the partition table and label to the disk.
  - h. Enter `q` to return to the `format` menu.
  - i. If you have finished creating slices, enter `q` to quit from the `format` utility. Otherwise, enter the `disk` command to select a new disk and repeat steps b to g to create or identify the slices on that disks.
4. On the other cluster nodes, use the `format` utility to determine whether the device names for the slices that you created are the same on every node.
  5. If you plan to use existing slices, enter the following command on every node to verify that they are not mounted as file systems:

```
df -k
```

This command displays information about the 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.

6. Enter commands similar to the following on every node to change the owner, group, and permissions on the character raw device file for each disk slice 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.

---

---

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

If the nodes are configured differently, the disk device names for the disks might be different on some nodes. Make sure that you specify the correct disk device name in these commands on each node.

---

---

7. If you also want to use raw devices for storage, refer to the "[Configure Raw Partitions or Raw Logical Volumes](#)" section on page 7-34.

Otherwise, refer to the "[Verify that the Cluster Software is Configured and Running](#)" section on page 7-44.

## Configure Raw Partitions or Raw Logical Volumes

Refer to the following section for information about configuring raw devices.

---

---

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

---

---

## Configuring Raw Partitions or Raw Logical Volumes

The following subsections describe how to configure raw partitions or raw logical volumes on Solaris:

- [Configure Raw Partitions or Raw Logical Volumes](#) on page 7-34  
Follow the procedures described in this section if you are installing RAC on a cluster without a supported cluster volume manager or a supported cluster file system.
- [SPARC Only: Configuring Raw Logical Volumes for Oracle CRS or Database File Storage](#) on page 7-39  
Follow the procedures described in this section only if you are using VERITAS Cluster Volume Manager (CVM) with Sun Cluster on SPARC systems.

### Configuring Raw Partitions for Oracle CRS or Database File Storage

This section describes how to configure shared raw partitions for Oracle CRS files (Oracle Cluster Registry and CRS voting disk) and Oracle database files.

A partition is a slice defined on a shared disk or on a shared disk array volume.

---

---

**Note:** For RAC installations, Solaris Volume Manager volumes or partitions are not currently supported for Oracle CRS or database file storage.

---

---



Table 7–1 lists the number and size of the raw disk devices that you must configure for database files. Table 7–2 lists the number and size of the raw disk devices that you must configure for CRS files.

---

**Note:** Solaris imposes limits on the number of slices that you can create on a single disk device. If you want to use slices, you might need to create them on more than one disk or volume.

---

**Table 7–1 Raw Partitions Required for Database Files on Solaris**

| Number                  | Size (MB)                         | Purpose and Sample Symbolic Link Name                                                                                                                             |
|-------------------------|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1                       | 500                               | SYSTEM tablespace:<br><i>dbname_system_raw_500m</i>                                                                                                               |
| 1                       | 300 + (Number of instances * 250) | SYSAUX tablespace:<br><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):<br><i>dbname_undotbs<math>n</math>_raw_500m</i>               |
| 1                       | 250                               | TEMP tablespace:<br><i>dbname_temp_raw_250m</i>                                                                                                                   |
| 1                       | 160                               | EXAMPLE tablespace:<br><i>dbname_example_raw_160m</i>                                                                                                             |
| 1                       | 120                               | USERS tablespace:<br><i>dbname_users_raw_120m</i>                                                                                                                 |
| 2 * number of instances | 120                               | Two online redo log files for each instance ( $n$ is the number of the instance and $m$ is the log number, 1 or 2):<br><i>dbname_redon<math>m</math>_raw_120m</i> |
| 2                       | 110                               | First and second control files:<br><i>dbname_control{1 2}_raw_110m</i>                                                                                            |
| 1                       | 5                                 | Server parameter file (SPFILE):<br><i>dbname_spfile_raw_5m</i>                                                                                                    |
| 1                       | 5                                 | Password file:<br><i>dbname_pwdfile_raw_5m</i>                                                                                                                    |

**Table 7–2 Raw Partitions Required for CRS Files on Solaris**

| Number | Size (MB) | Purpose and Sample Symbolic Link Name                                                                                                                                                                                                        |
|--------|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1      | 100       | Oracle Cluster Registry:<br><i>ora_ocr_raw_100m</i><br><b>Note:</b> You need to create this raw device only once on the cluster. If you create more than one database on the cluster, they all share the same Oracle Cluster Registry (OCR). |
| 1      | 20        | Oracle CRS voting disk:<br><i>ora_vote_raw_20m</i><br><b>Note:</b> You need to create this raw device 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 configure shared raw partitions for CRS files, database files, or both, follow these steps:

1. If you intend to use raw partitions for database file storage, 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. If necessary, install or configure the shared disks that you intend to use and restart the system.
3. If you want to use disk slices, follow these steps to create or identify the required disk slices:

- a. To list the disks attached to the system, 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`).

- b. Enter the number corresponding to the disk that you want to use.

---

---

**Note:** Make sure that the disk you choose is not being used for another purpose. For example, make sure that it is not being used as a component for a logical volume manager volume.

---

---

- c. On x86 systems only, use the `fdisk` command to create a Solaris partition on the disk if one does not already exist.

Solaris `fdisk` partitions must start at cylinder 1, not cylinder 0. If you create a new `fdisk` partition, you must label the disk before continuing.

- d. Enter the `partition` command, followed by the `print` command to display the partition table for the disk that you want to use.

- e. Identify or create slices for each of the partitions that you require.

The previous table shows the number and size of the partitions that you require for both CRS files and database files.

---

---

**Note:** To prevent the CRS or database files from overwriting the partition table, do not use slices that start at cylinder 0 (for example, slice 2).

---

---

- f. Make a note of the number of the slices that you want to use.

- g. If you modified a partition table or created a new one, enter the `label` command to write the partition table and label to the disk.

- h. Enter `q` to return to the `format` menu.

- i. After you have finished creating slices, enter `q` to quit from the `format` utility.
4. If you plan to use existing partitions, enter the following command on every node to verify that they are not mounted as file systems:

```
df -k
```

This command displays information about the 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.

5. On the other cluster nodes, use the `format` utility to determine whether the device names for the slices that you created are the same on every node.
6. If the device name for a slice that you want to use is different on any node, enter a command similar to the following on all nodes to create symbolic links to the appropriate device using a common name:

---



---

**Note:** If you are using a multi-pathing disk driver, make sure that you create the links to the correct logical device name for the slice.

The previous table provides suggested symbolic link names for each partition.

---



---

```
ln -s /dev/rdisk/cxydzsn /dev/rdisk/ora_ocr_raw_100m
```

7. Enter commands similar to the following on every node to change the owner, group, and permissions on the character raw device file for each partition:

- OCR:

```
chown root:oinstall /dev/rdisk/cxydzsn
chmod 640 /dev/rdisk/cxydzsn
```

- CRS voting disk and database files:

```
chown oracle:dba /dev/rdisk/cxydzsn
chmod 660 /dev/rdisk/cxydzsn
```

---



---

**Note:** If you are using a multi-pathing disk driver, make sure that you set the permissions only on the correct logical device name for the partition.

If you created symbolic links, specify the symbolic link paths in these commands.

---



---

8. If you are using raw disk devices for database files, follow these steps to create the DBCA raw device mapping file:

---



---

**Note:** You must complete this procedure only if you are using raw devices for database files. The DBCA raw device mapping file enables the DBCA to identify the appropriate raw disk device for each database file. You do not specify the raw devices for the Oracle CRS files in the DBCA raw device mapping file.

---



---

- a. 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
```
- b. 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.

- c. Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
- d. Using any text editor, create a text file similar to the following that identifies the device file name associated with each database file.

Oracle recommends that you use a file name similar to `dbname_raw.conf` for this file.

---

**Note:** The following example shows a sample mapping file for a two-instance RAC cluster. Some of the partitions use alternative symbolic link names. Make sure that the partition device file name that you specify identifies the same partition on all nodes.

---

```
system=/dev/rds/c2t1d1s3
sysaux=/dev/rds/c2t1d1s4
example=/dev/rds/c2t1d1s5
users=/dev/rds/c2t1d1s6
temp=/dev/rds/c2t1d2s3
undotbs1=/dev/rds/c2t1d2s4
undotbs2=/dev/rds/c2t1d2s5
redo1_1=/dev/rds/c2t1d2s6
redo1_2=/dev/rds/c2t1d3s3
redo2_1=/dev/rds/c2t1d3s4
redo2_2=/dev/rds/c2t1d3s5
control1=/dev/rds/c2t1d4s3
control2=/dev/rds/c2t1d4s3
spfile=/dev/rds/dbname_spfile_raw_5m
pwdfile=/dev/rds/dbname_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=device_file_name
```

The symbolic link names suggested in the previous table include the database object identifier that you must use in this mapping file. For example,

in the following symbolic link name, `redo1_1` is the database object identifier:

```
rac_redo1_1_raw_120m
```

- For a RAC database, the file must specify one automatic undo tablespace datafile (`undotbsn`) and two redo log files (`redo1_1`, `redo2_1`) 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 tablespace datafiles.
- e. Save the file and note the file name that you specified.
  - f. 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.
9. When you are installing Oracle Cluster Ready Services, you must enter the paths to the appropriate device files when prompted for the path of the OCR and CRS voting disk, for example:

```
/dev/rdsk/cxydzsn
```

10. To continue, refer to the "[Verify that the Cluster Software is Configured and Running](#)" section on page 7-44.

### SPARC Only: Configuring Raw Logical Volumes for Oracle CRS or Database File Storage

This section describes how to configure raw logical volumes using VERITAS Cluster Volume Manager (CVM) with Sun Cluster 3.1 on SPARC systems.

---

**Note:** For information about other supported options for creating logical volumes for Oracle CRS and database files, refer to the Certify page on the *OracleMetalink* Web site:

```
http://metalink.oracle.com/
```

---

### Create a Shared Disk Group

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 restart 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 on every node 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 (VxVM) disks. If you use a different LVM, refer to 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.

6. 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
```

7. 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:

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`.

- Identify the logical volumes that you must create.

[Table 7–3](#) lists the number and size of the logical volumes that you must create for database files. [Table 7–4](#) lists the number and size of the logical volumes that you must create for CRS files.

**Table 7–3 Raw Logical Volumes Required for Database Files on Solaris**

| Number                        | Size (MB)                                  | Purpose and Sample Logical Volume Name                                                                                                                                       |
|-------------------------------|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1                             | 500                                        | SYSTEM tablespace:<br><code>dbname_system_raw_500m</code>                                                                                                                    |
| 1                             | 300 +<br>(Number of<br>instances *<br>250) | SYSAUX tablespace:<br><code>dbname_sysaux_raw_800m</code>                                                                                                                    |
| Number<br>of<br>instances     | 500                                        | UNDOTBS $n$ tablespace (One tablespace for each instance, where $n$ is the number of the instance):<br><code>dbname_undotbs<math>n</math>_raw_500m</code>                    |
| 1                             | 250                                        | TEMP tablespace:<br><code>dbname_temp_raw_250m</code>                                                                                                                        |
| 1                             | 160                                        | EXAMPLE tablespace:<br><code>dbname_example_raw_160m</code>                                                                                                                  |
| 1                             | 120                                        | USERS tablespace:<br><code>dbname_users_raw_120m</code>                                                                                                                      |
| 2 *<br>number of<br>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):<br><code>dbname_redon<math>m</math>_raw_120m</code> |
| 2                             | 110                                        | First and second control files:<br><code>dbname_control{1 2}_raw_110m</code>                                                                                                 |
| 1                             | 5                                          | Server parameter file (SPFILE):<br><code>dbname_spfile_raw_5m</code>                                                                                                         |
| 1                             | 5                                          | Password file:<br><code>dbname_pwdfile_raw_5m</code>                                                                                                                         |

**Table 7–4 Raw Logical Volumes Required for CRS Files on Solaris**

| Number | Size (MB) | Purpose and Sample Logical Volume Name                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1      | 100       | Oracle Cluster Registry:<br><code>ora_ocr_raw_100m</code><br><br><b>Note:</b> 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 (OCR).<br><br>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. |

**Table 7–4 (Cont.) Raw Logical Volumes Required for CRS Files on Solaris**

| Number | Size (MB) | Purpose and Sample Logical Volume Name                                                                                                                                                                                                |
|--------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1      | 20        | Oracle CRS voting disk:<br>ora_vote_raw_20m<br><b>Note:</b> 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 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=oinstall mode=640
```

In this example, *diskgroup* is the name of the disk group you created previously, for example, *oradg*.

4. To create the other required logical volumes, enter commands 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 that 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 chose 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, group, 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_raw_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 disk 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 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 tablespace datafiles.
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:** This section applies to SPARC systems only.

Sun Cluster is not required. However, if it is installed, Oracle CRS can integrate with it.

---

---

To verify that the Sun Cluster software is configured and running on Solaris SPARC systems:

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, refer to 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.

---



---

**Note:** If no Oracle Net listeners are running, refer to the ["Configure the oracle User's Environment"](#) section on page 7-46 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 (`csh` or `tcsh`):

```
% vi .login
```

7. Enter or edit the following line, specifying a value of 022 for the default file mode 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 that 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 that 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 or identified earlier.

14. If you are using raw devices for database file 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:

- Bourne, Bash, or Korn shell:

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

- C shell:

```
% unsetenv ORACLE_HOME
% unsetenv TNS_ADMIN
```

---

---

**Note:** If the ORACLE\_HOME environment variable is set, the Installer uses the value that 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 that the environment variables you set in this section have the correct values.

---



---

## Pre-Installation Tasks for RAC on Windows

This chapter describes the tasks that you must complete before you start the Oracle Universal Installer to install RAC on Windows-based systems. The chapter includes following topics:

- [Oracle Database System Requirements](#)
- [Hardware and Software Certification](#)
- [Network Requirements](#)
- [Individual Component Requirements](#)

**See Also:** *Oracle Real Application Clusters Quick Installation Guide for Oracle Database Standard Edition for Windows for Oracle Database Standard Edition for Windows* for step-by-step installation instructions of RAC Standard Edition on a two-node cluster with no Oracle software previously installed.

### Oracle Database System Requirements

The following sections list the system requirements for Oracle Database 10g. Some individual components also have requirements that must be satisfied before installation.

- [Software Requirements](#)
- [Hardware Requirements](#)
- [Verifying Hardware Requirements](#)

### Software Requirements

[Table 8–1](#) lists the software requirements for Oracle Database 10g Real Application Clusters.

**Table 8–1** *Software Requirements*

| Requirement         | Value            |
|---------------------|------------------|
| System Architecture | 32-bit or 64-bit |

Note: Oracle provides both 32-bit and 64-bit versions of Oracle Database for Windows. Currently, the 64-bit version of the database must run on the 64-bit version of the operating system. The 32-bit version of the database must run on the 32-bit version of the operating system.

**Table 8–1 (Cont.) Software Requirements**

| Requirement      | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operating System | <ul style="list-style-type: none"> <li>■ RAC for Windows 32-bit is supported on the following operating systems: <ul style="list-style-type: none"> <li>– Windows 2000 with service pack 1 or higher. All editions, including Terminal Services and Windows 2000 MultiLanguage Edition (MLE), are supported.</li> <li>– Windows Server 2003</li> </ul> </li> <li>■ RAC for Windows 64-bit is supported on the following operating systems: <ul style="list-style-type: none"> <li>– Windows Server 2003 Datacenter Edition for 64-bit Itanium 2 Systems</li> <li>– Windows Server 2003 Enterprise Edition for 64-bit Itanium 2 Systems</li> </ul> </li> </ul> <p>Windows Multilingual User Interface Pack is supported on Windows Server 2003.</p>                                                                                                                                                                                                                                                                                                                                                                                              |
| Compiler         | <p>Windows 32-bit compilers:</p> <ul style="list-style-type: none"> <li>■ ACUCOBOL-GT for Pro*COBOL compiler version 6.0.0 is supported.</li> <li>■ Pro*COBOL supports the Micro Focus Net Express compiler. Object Oriented COBOL (OOCOBOL) specifications are not supported.</li> <li>■ The following components are not supported with the Microsoft Visual C++ 6.0, Microsoft Visual C++ .NET 2002, and Microsoft Visual C++ .NET 2003 compilers: <ul style="list-style-type: none"> <li>– Oracle C++ Call Interface</li> <li>– Oracle Call Interface</li> <li>– External callouts</li> <li>– PL/SQL native compilation</li> <li>– XDK</li> </ul> </li> </ul> <p>Windows 64-bit compilers:</p> <ul style="list-style-type: none"> <li>■ The following components are supported with the Microsoft Platform SDK and Intel Electron C 7.1 compilers: <ul style="list-style-type: none"> <li>– Oracle C++ Call Interface</li> <li>– Oracle Call Interface</li> <li>– External callouts</li> <li>– PL/SQL native compilation</li> <li>– XDK</li> </ul> </li> <li>■ Object Oriented COBOL (OOCOBOL) specifications are not supported.</li> </ul> |
| Network Protocol | The Oracle Net foundation layer for RAC supports TCP/IP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

If you are currently running a cluster version of Oracle on an operating system version that is no longer supported, such as Windows NT Server 4.0, then you must first upgrade your operating system before upgrading to Oracle Database 10g Real Application Clusters.

**See Also:**

- ["Telnet and Terminal Services Support"](#) on page 8-6



## Hardware Requirements

The following hardware components are required for Oracle Database:

- RAM: 512 MB for 32-bit systems, 1GB for 64-bit systems
- Virtual memory: double the amount of RAM on 32-bit systems, 512 MB for 64-bit systems
- Hard disk space: See [Table 8-2](#)
- Temp disk space: 100 MB for 32-bit systems, 140 MB for 64-bit systems
- Video adapter: 256 color
- Processor: 400 MHz minimum for Windows 32-bit, Itanium 2 or higher for Windows 64-bit

### See Also:

- ["Oracle Transparent Gateway Requirements"](#) on page 8-22 for additional requirements
- ["Configuring Disk Storage for Oracle Datafiles and Recovery Files"](#) on page 8-9
- ["Creating Directories for Oracle Datafiles or Recovery Files"](#) on page 8-10
- ["Configuring Disks for Automatic Storage Management"](#) on page 8-12
- ["Configuring Raw Partitions"](#) on page 8-18
- ["Installation on Windows-Based Systems with the Minimum Memory Requirements"](#) on page 11-12

## Hard Disk Space Requirements

This section lists space requirements for both the Enterprise and Standard Editions of Oracle Database 10g RAC. Oracle recommends installing Oracle components on Oracle Cluster File Systems (OCFS).

The space requirements listed in this section are more accurate than the hard disk values reported by the Oracle Universal Installer Summary screen. The Summary screen does not include accurate values for disk space, the space required to create a database (over 700 MB), or the size of compressed files that are expanded on the hard drive.

The hard disk requirements for Oracle Database components include 32 MB required to install Java Runtime Environment (JRE) and Oracle Universal Installer on the partition where the operating system is installed. If sufficient space is not detected, installation fails and an error message appears. [Table 8-2](#) lists the hard disk space requirements, including space for the starter database.

**Table 8-2** *Hard Disk Space Requirements*

| Operating System Version | System Drive | Oracle Home Drive |
|--------------------------|--------------|-------------------|
| 32-bit system            | 100 MB       | 1.5 GB            |
| 64-bit system            | 140 MB       | 3 GB              |

## Oracle Cluster File System Pre-Installation Steps

Perform the pre-installation steps described in this section to setup up the shared disk raw partitions for OCFS. Windows refers to raw partitions as logical drives. If you need more information about creating partitions, refer to the Windows online help from within the disk administration tools.

1. Run Windows Disk Management from one node to create an extended partition. Use a basic disk: dynamic disks are not supported.

Note that the Globally Unique Identifier (GUID) Partition Type (GPT), available on 64-bit systems, is not supported for OCFS. You must use the Master Boot Record (MBR) partition style.

2. Create at least two partitions: one for the Oracle home and one for the Oracle database files.

You do not need to create a partition for the voting disk if you plan to use CFS. CFS stores the voting device in a file.

The number of partitions used for CFS affects performance. Therefore, you should create the minimum number of partitions needed for the CFS option you choose.

To create the required partitions, perform the following steps:

1. From one of the existing nodes of the cluster, run the Windows disk administration tool as follows:
  - Navigate to **Start > Settings > Control Panel > Administrative Tools > Computer Management**
  - Expand the Storage folder to Disk Management. Use a basic disk with a Master Boot Record (MBR) partition style as an extended partition for creating partitions.
2. Right click inside an unallocated part of an extended partition and choose Create Logical Drive. A wizard presents pages for configuring the logical drive. Select the select logical drive radio button and click Next.
3. Enter the size that you want for the partition and click Next.
4. Choose the option "Do not assign a drive letter or path", click Next, and then choose the option "Do not format this partition". Click Finish on the last page of the wizard.
5. Repeat steps 2 through 4 for the second and any additional partitions. An optimal configuration is one partition for the Oracle home and one partition for Oracle database files.
6. If you are preparing drives on a Windows 2003 or on a 64-bit system, then you should restart all nodes in the cluster after you have created the logical drives.
7. Check all nodes in the cluster to ensure that the partitions are visible on all the nodes and to ensure that none of the Oracle partitions have drive letters assigned. If any partitions have drive letters assigned, then remove them by performing these steps:
  - Right-click the partition in the Windows disk administration tool
  - Select "Change Drive Letters and Paths..." from the menu
  - Click **Remove** in the "Change Drive Letter and Paths" window

## Verifying Hardware Requirements

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

1. Determine the physical RAM size. For a computer using Windows 2000, for example, open **System** in the control panel and select the **General** tab. If the size of the physical RAM installed in the system is less than the required size, then you must install more memory before continuing.
2. Determine the size of the configured swap space (also known as paging file size). For a computer using Windows 2000, for example, open **System** in the control panel, select the **Advanced** tab, and click **Performance Options**.

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

3. Determine the amount of free disk space on the system. For a computer using Windows 2000, for example, open **My Computer**, right-click the drive where the Oracle software is to be installed, and choose **Properties**.
4. Determine the amount of disk space available in the `temp` directory. This is equivalent to the total amount of free disk space, minus what will be needed for the Oracle software to be installed.

On Windows 32-bit systems, you require 100 MB of disk space available in the `temp` directory, and on Windows 64-bit systems, you require 140 MB of temp disk space. If you do not have sufficient space, then first delete all unnecessary files. If the temp disk space is still less than the required amount, then set the `TEMP` or `TMP` environment variable to point to a different hard drive. For a computer using Windows 2000, for example, open the **System** control panel, select the **Advanced** tab, and click **Environment Variables**.

## Hardware and Software Certification

The platform-specific hardware and software requirements included in this installation guide were current at the time this guide was published. However, because new platforms and operating system software versions might be certified after this guide is published, review the certification matrix on the Oracle*MetaLink* Web site for the most up-to-date list of certified hardware platforms and operating system versions. This Web site also provides compatible client and database versions, patches, and workaround information for bugs. The Oracle*MetaLink* Web site is available at the following URL:

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

You must register online before using Oracle*MetaLink*. After logging in, select **Certify & Availability** from the left-hand column. From the Product Lifecycle page, select the **Certifications** button. Other Product Lifecycle options include **Product Availability**, **Desupport Notices**, and **Alerts**.

The following sections list the following certification information:

- [Web Browser Support](#)
- [Telnet and Terminal Services Support](#)

### Web Browser Support

The following Web browsers are supported for *iSQL\*Plus* and Oracle Enterprise Manager Database Control:

- Netscape Navigator 7.0 and higher
- Microsoft Internet Explorer 5.5, 6.0, and higher with service pack 1
- Mozilla version 1.3.1 and higher
- On 64-bit systems only, Microsoft Internet Explorer 6.0 with service pack 2 Web browser

## Telnet and Terminal Services Support

This section contains these topics:

- [Windows Telnet Services Support](#)
- [Windows Terminal Services and Remote Desktop Support](#)

### Windows Telnet Services Support

Windows 2000 and Windows Server 2003 include a Telnet Service that enables remote users to log on to the operating system and run console programs using the command line. Oracle supports the use of database command line utilities such as `sqlplus`, `export`, `import` and `sqlldr` using this feature, but does not support the database GUI tools such as Oracle Universal Installer, Database Configuration Assistant, and Oracle Net Configuration Assistant.

---

---

**Note:** Ensure that the Telnet service is started on the Services control panel.

---

---

### Windows Terminal Services and Remote Desktop Support

Oracle Database 10g Real Application Clusters supports Terminal Services on Windows 2000 and Windows Server 2003. However, installation of Oracle components from a remote Terminal Services Client on to a Windows server that is running a Terminal Server service is only supported for 32-bit systems. If you encounter problems with the installation through Terminal Server on 32-bit systems, then Oracle recommends connecting to the Terminal Services console session of the server (using `mstsc\console`). On 64-bit systems, start all configuration tools from the Terminal Server console and not from the Terminal Services Client.

Platform-specific support information is as follows:

- Windows 2000: Oracle supports installing, configuring, and running Oracle Database from a remote Terminal Services Client.
- Windows Server 2003: You can configure Windows Server 2003 to use Terminal Services in Remote Desktop for Administration Mode or Terminal Server Mode.

The following products and features are not supported with Windows Terminal Services:

- Oracle Connection Manager
- Oracle Object Link Manager
- Oracle Services for Microsoft Transaction Server
- Server Management (SRVM)

**See Also:**

- The Microsoft Web site for more information about terminal servers

<http://www.microsoft.com/>

- The Oracle*MetaLink* Web site for the latest Terminal Server certification information

<http://metalink.oracle.com>

## 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, refer to 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 private and public network interface names must be different from each other.
- Network interfaces names may not contain multi-byte or unicode characters or the character . (dot or period).
- Interface names are case sensitive. If the interface names are not identical with respect to case, then the VIP functionality can fail.
- The private network interface name must be the same on all nodes.
- The public network interface name must be the same on all nodes.
- The public interface must be listed first in the ipconfig list.
- Oracle supports the TCP/IP protocol for the public and private networks.

## 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 `%SystemRoot%\system32\drivers\etc\hosts` file on each node to associate private host names with private IP addresses.

For example, in a two node cluster where each node has one public and one private interface, you might have the following host names and IP addresses:

| Host Name             | Type    | IP Address    | Registered In                           |
|-----------------------|---------|---------------|-----------------------------------------|
| rac1.mydomain.com     | Public  | 143.46.43.100 | DNS                                     |
| rac2.mydomain.com     | Public  | 143.46.43.101 | DNS                                     |
| rac1-vip.mydomain.com | Virtual | 143.46.43.104 | DNS                                     |
| rac2-vip.mydomain.com | Virtual | 143.46.43.105 | DNS                                     |
| rac1-priv             | Private | 10.0.0.1      | %SystemRoot%\system32\drivers\etc\hosts |
| rac2-priv             | Private | 10.0.0.2      | %SystemRoot%\system32\drivers\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 `%SystemRoot%\system32\drivers\etc\hosts` file on all nodes, specifying the private IP address and associated private host name:

```
10.0.0.1 rac1-priv
```

If you need to change a network interface name, follow these steps:

1. Depending on your system architecture, complete the first step as follows:
  - On a Windows 32-bit system, navigate to **Start > Settings > Control Panel > Network and Dial-up Connections**
  - On a Windows 64-bit system, navigate to **Start > Settings > Control Panel > Network Connections**
2. Right click the icon of the network interface for which you need to change the name
3. Select Rename
4. Enter and save the new name

To ensure that your public interface appears first in your ipconfig list, follow these steps:

1. Depending on your system architecture, complete the first step as follows:
  - On a Windows 32-bit system, navigate to **Start > Settings > Control Panel > Network and Dial-up Connections**
  - On a Windows 64-bit system, navigate to **Start > Settings > Control Panel > Network Connections**

2. In the **Advanced** menu, click **Advanced Settings...**
3. If the public interface name is not the first name listed under the **Adapters and Bindings** tab, then select it and click the arrow to move it to the top of list
4. Click **OK** to save the setting and then exit network setup dialog

## Individual Component Requirements

This section contains these topics:

- [Configuring Disk Storage for Oracle Datafiles and Recovery Files](#)
- [Creating Directories for Oracle Datafiles or Recovery Files](#)
- [Configuring Disks for Automatic Storage Management](#)
- [Configuring Raw Partitions](#)
- [Oracle Advanced Security Requirements](#)
- [Oracle Enterprise Manager Requirements](#)
- [Oracle Managed Files Requirements](#)
- [Oracle Transparent Gateway Requirements](#)

### Configuring Disk Storage for Oracle Datafiles and Recovery Files

As a first step in configuring your disk storage for RAC, you must disable operating system write caching on all disks that will be used to share data between nodes in your cluster. To disable operating system write caching, perform these steps:

1. Navigate to **Start > Settings > Control Panel > Administrative Tools > Computer Management > Device Manager > Disk drives**
2. Expand the **Disk drives** hive and double-click the first drive listed
3. Under the Disk Properties tab for the selected drive, uncheck the option that enables the write cache
4. Double-click each of the other drives listed in the **Disk drives** hive and disable the write cache as described in the previous step

---



---

**Caution:** Any disks that you use to store files, including database files, that will be shared between nodes, must have operating system write caching disabled. You should not disable write caching provided by a storage vendor to improve storage performance.

---



---

On Windows 2003 only, a second requirement is to enable the disk automount feature by performing the following steps on all nodes of your cluster:

1. Enter the command `diskpart` in a Command window
2. Enable the disk automount feature by entering the `automount enable` command and confirming its successful execution as follows:

```
DISKPART> automount enable
Automatic mounting of new volumes enabled.
```

3. Type `exit` to terminate the Diskpart session

When you have prepared all the cluster nodes in your Windows 2003 system as described in the previous steps, restart all of the nodes.

The remainder of this section describes the storage options for storing Oracle datafiles and optionally, Oracle database recovery files. After you choose the storage method that you want to use for each file type, refer to 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 except for Oracle Real Application Clusters Standard Edition installations. For Standard Edition databases, you must use Automatic Storage Management (ASM) for your database file storage.

---

---

### Choosing a Storage Option for Oracle Datafiles

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

- Cluster file system
- Automatic Storage Management
- Raw Devices

### 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
- 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 datafiles.

### Configuring Disk Storage

For more information about these options, refer to the "[Storage Considerations for Installing Oracle Database 10g Real Application Clusters](#)" section on page 1-6. For information about how to configure disk storage before you start the installation, refer to one of the following sections depending on your choice:

- To use a cluster file system for database or recovery file storage, refer to the "[Creating Directories for Oracle Datafiles or Recovery Files](#)" section on page 8-10.
- To use ASM for database or recovery file storage, refer to the "[Configuring Disks for Automatic Storage Management](#)" section on page 8-12.
- To use raw devices for database file storage, refer to the "[Configuring Raw Partitions](#)" section on page 8-18.

## Creating Directories for Oracle Datafiles or Recovery Files

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

### Guidelines for Placing Oracle Datafiles on a File System

If you decide to place the Oracle datafiles on a cluster file system, use the following guidelines when deciding where to place them:



- You can choose either a single cluster file system or more than one cluster file system to store the datafiles:
  - If you want to use a single cluster file system, choose a cluster 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 cluster file system, choose cluster 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 cluster file system (or systems) that you choose must have at least 3 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 cluster file systems that you choose should be on physical devices that are used only by the database.
- The default location suggested by Oracle Universal Installer for the database file directory is a subdirectory of the Oracle base directory. However, this default location is not appropriate for RAC production databases.

### 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 cluster file system, use the following guidelines when deciding where to place them:

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

---



---

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

---



---

- The cluster file system that you choose should have at least 3 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 to specify a different value.

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

- The default location suggested by Oracle Universal Installer for the recovery area directory is a subdirectory of the Oracle base directory. However, this default location is not appropriate for RAC production databases.

## Creating Required Directories

---



---

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

---



---

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

1. Use Windows Explorer to determine the free disk space on the file system.
2. From the display, identify the file systems that you want to use:

| File Type      | File System Requirements                                                                                                                                                                                                         |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Datafiles      | Choose either: <ul style="list-style-type: none"> <li>■ A single cluster file system with at least 3 GB of free disk space</li> <li>■ Two or more cluster file systems with at least 3 GB of free disk space in total</li> </ul> |
| Recovery files | Choose a cluster file system with at least 3 GB of free disk space.                                                                                                                                                              |

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

3. Note the names of the directories for the file systems that you identified.
4. If you also want to use ASM or raw devices for storage, refer to one of the following sections:
  - [Configuring Disks for Automatic Storage Management](#)
  - [Configuring Raw Partitions](#)

Otherwise refer to the [Stop Existing Oracle Services](#) section on page 8-21.

## Configuring 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](#)

## 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 for your ASM disk groups. In addition to the ASM files, CRS requires two key files that must be located in logical drives on the shared disks: one for a voting disk that requires 20 MB of storage, and one for the Oracle Cluster Registry that requires 100 MB of storage.

To configure the ASM disk space, follow these steps:

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

---



---

**Note:** You do not have to use the same storage mechanism for datafiles and recovery files. Datafiles can be placed on ASM while the recovery area can be on a file system.

---



---

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 Database Configuration Assistant 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 Database Configuration Assistant after the installation to create a database.

- If you select an installation type that runs Database Configuration Assistant in non-interactive mode, you must use the same ASM disk group for datafiles 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        |

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 Database Configuration Assistant 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, refer to the ["Configuring Disks for ASM"](#) section on page 8-16.

### Using an Existing ASM Disk Group

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

- If you select an installation method that runs Database Configuration Assistant 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 Database Configuration Assistant after the installation to create a database.

- If you select an installation type that runs Database Configuration Assistant in non-interactive mode, you should choose an existing disk group for the new database; you cannot create a new disk group.

---



---

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

---



---

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

1. Check the Services Control Panel for the OracleASMSERVICE+ASMn service, where n is the node number.
2. Temporarily 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 AS SYSDBA privilege and start the instance if necessary:

```
C:\>ORACLE_BASE\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

Automatic Storage Management (ASM) with RAC is supported on Windows 2000 and Windows Server 2003. You must use basic disks for your ASM storage, dynamic disks are not supported. To use ASM with direct attached storage (DAS) or storage area network (SAN) storage, the disks must be stamped with a header by `asmtool` or `asmtoolg` (GUI version).

In order to use a DAS or SAN disk in ASM, the disk must have a partition table. Oracle recommends creating exactly one partition for each disk containing the entire disk. Use Microsoft Computer Management or the command line tool `diskpart` to create the partition. Once the partitions have been created, unassign the drive letters and run `asmtoolg` or `asmtool`. These tools associate meaningful, persistent names with disks to facilitate using those disks with ASM. ASM uses disk strings to more easily operate on groups of disks at once, so the names created by `asmtool` make this easier than using Windows drive letters.

All disk names created by `asmtool` begin with the prefix `ORCLDISK` for identification purposes. They can be used as raw devices in the ASM instance by specifying a name `\\.\ORCLDISKn`.

**See Also:** ["Assigning Logical Names or Drive Letters, or Mounting Directories"](#) on page 8-18 for more information about using `diskpart` to create a partition

### Using `asmtoolg`

`asmtoolg` is a graphical interface for creating ASM device names. Use `asmtoolg` to add, change, delete, and examine the devices available for use in ASM.

---

---

**Note:** Use the `asmtool` command line to create an empty file. See the ["Using `asmtool` \(command line\)"](#) section on page 8-17.

---

---

To add or change disk stamps:

1. Double-click `asmtoolg`.
2. Select the **Add or change label** option, then click **Next**.

`asmtoolg` will show the devices available on the system. Unrecognized disks are labeled as "Candidate device", raw device files as "Oracle raw device file",

stamped ASM disks as "Stamped ASM disk", and unstamped ASM disks as "Unstamped ASM disks." The tool also shows disks that are recognized by Windows as a file system (such as NTFS). These are not available for use as ASM disks and cannot be selected. In addition Microsoft Dynamic disks are not available for use as ASM disks.

3. On the Stamp Disks screen, select the disks to stamp.

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

You can also specify the stamps of individual devices.

4. Optionally, select a disk to edit the individual stamp (ASM link name).
5. Click **Next**.
6. Click **Finish**.

To delete disk stamps:

1. Select the **Delete labels** option., then click **Next**.

The delete option is only available if disks exist with stamps. The delete screen shows all stamped ASM disks.

2. On the Delete Stamps screen, select the disks to unstamp.
3. Click **Next**.
4. Click **Finish**.

### Using `asmtool` (command line)

`asmtool` is a command-line interface for stamping disks. It has the following options:

| Option                  | Description                                                                                                                                                                                                                                                                  | Example                                                                                                                           |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <code>-add</code>       | Adds or changes stamps. You must specify the hard disk, partition, and new stamp name. If the disk is a raw device or has an existing ASM stamp, then you must specify the <code>-force</code> option.                                                                       | <code>asmtool -add [-force]<br/>\Device\Harddisk1\Partition1 ORCLDISKASM0<br/>\Device\Harddisk2\Partition1 ORCLDISKASM1...</code> |
| <code>-addprefix</code> | Adds or changes stamps using a common prefix to generate stamps automatically. The stamps are generated by concatenating a number with the prefix specified. If the disk is a raw device or has an existing ASM stamp, then you must specify the <code>-force</code> option. | <code>asmtool -addprefix ORCLDISKASM [-force]<br/>\Device\Harddisk1\Partition1<br/>\Device\Harddisk2\Partition1...</code>         |
| <code>-list</code>      | List available disks. The stamp, windows device name, and disk size in megabytes are shown. Some disks may be file systems, and cannot be stamped. If the disk is a raw device or has an existing ASM stamp, then you must specify the <code>-force</code> option.           | <code>asmtool -list [-force]</code>                                                                                               |
| <code>-delete</code>    | Removes existing stamps from disks.                                                                                                                                                                                                                                          | <code>asmtool -delete ORCLDISKASM0 ORCLDISKASM1...</code>                                                                         |
| <code>-create</code>    | Creates an empty file.                                                                                                                                                                                                                                                       | <code>asmtool -create path_to_file size_in_megabytes</code>                                                                       |

---

---

**Note:** For `-add`, `-addprefix`, and `-delete`, `asmtool` will notify any ASM instances on the local machine and other nodes in the cluster, if available, to rescan the available disks.

---

---

## Configuring Raw Partitions

This section contains the following topics:

- [Creating Partitions and Logical Drives](#)
- [Assigning Logical Names or Drive Letters, or Mounting Directories](#)
- [Creating Raw Logical Volumes in a New Disk Group](#)

### Creating Partitions and Logical Drives

To create and configure partitions, use the disk administration tools provided by the operating system or third party vendors. The following administration tools are provided by the operating system:

- Windows 2000 and Windows Server 2003 provide Disk Management snap-in  
To access this tool, type `diskmgmt.msc` at the command prompt or navigate to **Start > Programs > Administrative Tools > Computer Management**. Then select the **Disk Management** node in the Storage tree.
- Windows Server 2003 provides a command line tool to manage disks.  
To access this tool, type `diskpart.exe` at the command prompt.

---

---

**Note:** `Diskpart.exe` for Windows 2000 is part of Windows 2000 Resource Kit. To download this tool, consult Microsoft documentation on the Microsoft Web site

<http://www.microsoft.com/>

---

---

**See Also:** The online help or documentation for the administration tool you are using

### Assigning Logical Names or Drive Letters, or Mounting Directories

After creating volumes, assign logical names or drive letters, or mount them on directories for use by Oracle. Use the Windows graphical interface or the command line to create a mounted drive or assign a drive letter to the partition, logical drive or volume. Names can be assigned to partitions using Oracle Object Link Manager.

The following example uses the `diskpart` tool to create a 32 MB partition on disk 100 and to assign the drive letter B to the partition. In this example, `diskpart.exe` is the command line tool for managing disks.

```
c:\> diskpart.exe
diskpart> select disk 100
diskpart> create partition primary size=32
diskpart> assign letter=B
```

Optionally, mount the partition on an NTFS folder instead of assigning a drive letter:

```
diskpart> assign mount=C:\mnt\raw_1
```



---

**Note:** You cannot use partitions defined on disks with primary partitions, as shown in the preceding examples, for storage defined while running the Oracle Universal Installer to install Cluster Ready Services as described in [Chapter 10, "Installing Cluster Ready Services on Windows"](#).

---

### Creating Raw Logical Volumes in a New Disk Group

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

1. Choose a name for the database that you want to create.
2. Create the logical volumes listed in the following table. You must create these volumes in order to install an Oracle database.

| Number              | Partition Size (MB) | Purpose and Sample Logical Volume Name                                                                                                   |
|---------------------|---------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| 1                   | 500                 | SYSTEM tablespace: <i>dbname_system_raw_500m</i>                                                                                         |
| 1                   | 500                 | SYSAUX tablespace: <i>dbname_sysaux_raw_500m</i>                                                                                         |
| 1 for each instance | 500                 | UNDOTBS1 tablespace: <i>dbname_undotbs1_raw_500m</i>                                                                                     |
| 1                   | 160                 | EXAMPLE tablespace: <i>dbname_example_raw_160m</i>                                                                                       |
| 1                   | 120                 | USERS tablespace: <i>dbname_users_raw_120m</i>                                                                                           |
| 2 for each instance | 120                 | Two online redo log files (where <i>m</i> is the group number and <i>n</i> is the log number, 1 or 2):<br><i>dbname_redom_n_raw_120m</i> |
| 2                   | 110                 | First and second control files:<br><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>                                                                              |

3. To create the other required logical volumes, using the command-line interface, enter a command similar to the following:

```
c:\> diskpart.exe
DISKPART> select disk diskn
DISKPART> create partition primary size=sizen
DISKPART> assign mount=folder
```

In this example:

- `diskpart.exe` is the command line tool for managing disks
- `diskn` is the disk number where the partitions are created
- `sizen` is the size of the partition, for example 500 represents 500 Megabytes
- `folder` is the absolute path to the NTFS where the partitions is mounted.

The following example shows a sample command to create 500 MB partition on disk 5 for the SYSAUX tablespace of a database named test:

```
c:\> diskpart.exe
DISKPART> select disk 5
```

```
DISKPART> create partition primary size=500
DISKPART> assign mount=D:\oracle\product\10.1.0\oradata\test\test_sysaux_500m
```

## 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 symbolic links 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, as in this example:

```
C:\>set ORACLE_BASE = E:\oracle
```

2. Create a database subdirectory under the Oracle base directory as in this example:

```
C:\>mkdir E:\oracle\dbname
```

where *dbname* is the name of the database that you chose previously.

3. Change directory to the %ORACLE\_BASE%\dbname directory.
4. Using any text editor, create a file called `conf.txt`. The file should have the following characteristics:

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

```
database_object_identifier = symbolic link name
```
- For your RAC database, the file must specify one automatic undo tablespace datafile (`undotbsn`) and two redo log files (`redon_1`, `redon_2`) for each instance where *n* is the instance number.
- Specify at least two control files (`control1`, `control2`).

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

```
system=\\.\dbname_SYSTEM
sysaux=\\.\dbname_SYSAUX
spfile=\\.\dbname_SPFILE
users=\\.\dbname_USERS
temp=\\.\dbname_TEMP
undotbs1=\\.\dbname_UNDOTBS1
undotbs2=\\.\dbname_UNDOTBS2
control1=\\.\dbname_CONTROL1
control2=\\.\dbname_CONTROL2
redo1_1=\\.\dbname_REDO1_1
redo1_2=\\.\dbname_REDO1_2
redo2_1=\\.\dbname_REDO2_1
redo2_2=\\.\dbname_REDO2_2
example=\\.\dbname_EXAMPLE
```

---

---

**Note:** After you complete your CRS installation, run `CRS_home\bin\GUIOracleObjectManager.exe` to create persistent symbolic links to the corresponding raw partitions.

---

---

5. Save the file and note the file name that you specified.
6. You may optionally set an environment variable, `DBCA_RAW_CONFIG`, to specify the full path to this file. For the Oracle base defined in Step 1, you would use the following command:

```
C:\>set DBCA_RAW_CONFIG=E:\oracle\dbname\conf.txt
```

## Stop Existing Oracle Services

---



---

**Attention:** If you are installing additional Oracle Database components in an existing Oracle home, stop **all** processes running in the Oracle home.

---



---

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, Oracle Universal 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 Oracle Universal Installer.

To stop an Oracle service, perform one of the following procedures:

- Using the GUI:
  1. Navigate to **Start > Settings > Control Panel > Administrative Tools > Services**
  2. Right click the service you want to stop
  3. Click **Stop**
- Using command line, enter the command:

```
C:\>net stop service
```

where *service* is the name of the service you want to stop

## Oracle Advanced Security Requirements

Satisfy hardware and software requirements to use authentication support with Oracle components. Some Oracle Advanced Security components can use a Lightweight Directory Access Protocol (LDAP) directory such as Oracle Internet Directory (provided on the installation media). Enterprise User Security also requires an LDAP directory.

**See Also:** *Oracle Advanced Security Administrator's Guide*

## Oracle Enterprise Manager Requirements

All Oracle Enterprise Manager products must be of the same release. Older versions of Enterprise Manager are not supported with the current release.

---



---

**Note:** All Oracle Enterprise Manager products, except Oracle Enterprise Manager Database Control and the Enterprise Manager Java Console, are released on the Enterprise Manager Grid Control installation media. Enterprise Manager Database Control is available on the Oracle Database installation media and Enterprise Manager Java Console is available on the Oracle Client installation media.

---



---

**See Also:** *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* available on the Enterprise Manager Grid Control installation media

## Oracle Managed Files Requirements

If you use OCFS or ASM for your database files, then your database will be created with Oracle-managed files by default. You may also elect to use Oracle-managed files if you choose the Custom installation type or the Advanced database creation option. If you use this feature, you need only specify the database object name instead of file names when creating or deleting database files.

Configuration procedures are required in order to enable Oracle Managed Files.

**See Also:** "Using Oracle-Managed Files" in *Oracle Database Administrator's Guide*

## Oracle Transparent Gateway Requirements

Oracle transparent gateways are only supported on 32-bit systems. This section contains these gateway topics:

- [Oracle Transparent Gateway Supported Configurations](#)
- [Oracle Transparent Gateway System Requirements](#)

### Oracle Transparent Gateway Supported Configurations

Table 8–3 provides gateway configurations tested by Oracle at the time of this document release. Oracle continually updates supported gateway configurations. For the latest supported configuration information, visit the following Web site

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

After you log in, select **Certify & Availability** from the left-hand column. From the Product Lifecycle page, select the **Certifications** button.

**Table 8–3 Supported Software for Oracle Transparent Gateway Components**

| Oracle Transparent Gateway                          | Requirement                                                                                                                                                                                                 |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle Transparent Gateway for Microsoft SQL Server | Microsoft SQL Server Version 7.0 or SQL Server 2000<br>Microsoft SQL Server ODBC Drivers                                                                                                                    |
| Oracle Transparent Gateway for Sybase               | Sybase Server, version 12.0, or 12.5 is required. If Sybase Server is not on the same computer as the gateway, then the version of Sybase Open client library certified for your Sybase Server is required. |
| Oracle Transparent Gateway for Teradata             | Teradata V2R.03.00.02 or V2R.04.00.0115<br>NCR Teradata ODBC Driver version 02.08.00.00                                                                                                                     |

**Table 8–3 (Cont.) Supported Software for Oracle Transparent Gateway Components**

| Oracle Transparent Gateway              | Requirement                                                                                     |
|-----------------------------------------|-------------------------------------------------------------------------------------------------|
| Oracle Transparent Gateway for IBM DRDA | See <i>Oracle Transparent Gateway for DRDA Installation and User's Guide</i> for details.       |
| Oracle Procedural Gateway for APPC      | See <i>Oracle Procedural Gateway for APPC Installation and Configuration Guide</i> for details. |

### Oracle Transparent Gateway System Requirements

Table 8–4 lists the system requirements for Oracle Transparent Gateways for Microsoft SQL Server, Sybase, and Teradata. Each Oracle Transparent Gateway has the same memory, networking, and disk space requirements.

**See Also:** *Oracle Transparent Gateway for DRDA Installation and User's Guide* and *Oracle Procedural Gateway for APPC Installation and Configuration Guide* for system requirements for the IBM DRDA and APPC gateways

These guides are available on the Oracle Database installation media.

**Table 8–4 System Requirements for Oracle Transparent Gateways**

| Requirement       | Value                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk Space        | 200 MB of additional disk space.                                                                                                                                                                                                                                                                                                                                                                                             |
| Memory            | 26 MB of real memory is recommended to support the gateway. The total real memory requirement for the concurrent use of the gateway also depends on these factors: <ul style="list-style-type: none"> <li>■ The SQL statement issued by the user</li> <li>■ The number of cursors currently opened against Microsoft SQL, Sybase, or Teradata server</li> <li>■ The number of columns in the table being accessed</li> </ul> |
| Oracle Networking | Oracle Net Services and Oracle Protocol Support for Named Pipes or TCP/IP must be installed on the Oracle database computer and the gateway computer. The Oracle Net Services components are included on the Oracle Database installation media.                                                                                                                                                                             |



# 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 9, "Installing Cluster Ready Services on UNIX"](#)
- [Chapter 10, "Installing Cluster Ready Services on Windows"](#)
- [Chapter 11, "Installing Oracle Database 10g with Real Application Clusters"](#)
- [Chapter 12, "Creating RAC Databases with the Database Configuration Assistant"](#)
- [Chapter 13, "Real Application Clusters Post-Installation Procedures"](#)





---

---

# Installing Cluster Ready Services on UNIX

This chapter describes the procedures for installing Cluster Ready Services (CRS) on UNIX, phase one of the Oracle Database 10g Real Application Clusters installation on UNIX-based systems. The topics in this chapter are:

- [Installation Setup Procedures](#)
- [Installing the UDLM for Sun Clusters](#)
- [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 the installer executes `ssh` or `rsh` commands, including any login or other shell scripts they invoke, you may see errors about invalid arguments on standard input if the scripts generate any output. Although these are benign errors, you can avoid them by removing all commands that generate output when you execute `ssh` or `rsh` commands.

---

---

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 inter-node, or instance-to-instance Cache Fusion traffic. Using public interfaces for Cache Fusion can cause performance problems.

## Installing the UDLM for Sun Clusters

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 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.

---



---

### Notes:

- On Linux platforms, you cannot install the Oracle Database 10g Cluster Ready Services software on the Oracle Database 10g Linux cluster file system.
  - On Mac OS X systems, you must run the Installer from the terminal session where you logged in as the oracle user and set the user's environment (described in [Chapter 6, "Pre-installation Tasks for RAC on Mac OS X"](#)).
- 
- 

1. If a GSD from Oracle9i, Release 9.2 or earlier, is running, then stop it before installing Oracle Database 10g CRS by executing the following command from the GSD's Oracle home:

```
$ORACLE_HOME/bin/gsdctl stop
```

If you are installing CRS on a node that already has a single-instance Oracle Database 10g installation, then perform the following operations:

- a. If there are any ASM instances running on the node, stop these existing ASM instances. After CRS is installed, restart the ASM instances again. This ensures that they are associated with the new cluster software.
- b. Unless you are using a Mac OS X-based system, log on as the root user and stop all Oracle services by executing the command:

```
Oracle home/bin/localconfig delete
```

where *Oracle home* is the home that is running CSS.

If you are using a Mac OS X-based system, then execute the following command as a member of the admin group to stop all Oracle services:

```
sudo Oracle home/bin/localconfig delete
```

where *Oracle home* is the home that is running CSS.

2. Log in as the `oracle` user and set the `ORACLE_BASE` environment variable to specify the Oracle base directory that you created previously, for example `/u01/app/oracle`.
3. Set the `ORACLE_HOME` environment variable to specify the CRS home directory that you created previously, for example `/u01/crs/oracle/product/10.1.0/crs_1`.
4. Run the `runInstaller` command from the top-level directory of the Oracle Cluster Ready Services Release 1 CD-ROM or the `crs` directory on the DVD-ROM. These are separate CD-ROMs and DVD-ROMs that contain the Cluster Ready Services software. When the OUI displays the Welcome page, click Next.
5. 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 7.
  - 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.
6. The OUI dialog indicates that you should run the `oraInventory location/orainstRoot.sh` script. Run the `orainstRoot.sh` script as `root` user, or with `sudo` command prefix if you are using a Mac OS X system. When the script completes, click Continue and the OUI displays the Specify File Locations page.
7. 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.

---

---

8. In the Language Selection page select the languages that you want CRS to use, click Next, and the OUI displays the Cluster Configuration page.
9. The Cluster Configuration page contains pre-defined node information if the OUI detects that your system has vendor clusterware. Otherwise, the OUI displays the Cluster Configuration 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 host names, that is, underscores (`_`), hyphens (`-`), and single-byte alphanumeric characters (a to z, A to Z, and 0 to 9). 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 Specify Network Interface Usage page.

---

---

**Note:** The IP addresses that you use for all of the nodes in the current installation process must be from the same subnet.

---

---

10. In the Specify Network Interface Usage 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` unless you are installing on a Sun Cluster with version 3.1, Update 1, or earlier. If you are installing CRS on a Sun Cluster up to version 3.1 Update 1, then do not change any of the interfaces from the default value of `Do Not Use`. This enables Sun Cluster make the correct choice. For versions of Sun Cluster after 3.1 Update 1, you should denote the special cluster interface as `Private` on the Private Interconnect Enforcement page.
11. When you click Next on the Specify Network Interface Usage page, the OUI will look for the `ocr.loc` file. The OUI will look for the file in the `/etc/oracle` directory in AIX, Mac OS X, and Linux environments. On other UNIX systems, the OUI will look for the `ocr.loc` file in the `/var/opt/oracle` directory. If the `ocr.loc` file exists, 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 12.  
  
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.
12. 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

13. 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.

14. The OUI displays the list of components that will be installed on the Summary page. Review the list of components and click install.

During the installation, unless you are using a clustered file system for the CRS home, 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.

15. Before running the `root.sh` script, log in as `root` and change permissions on the parent directories of the CRS home directory to permit only the `root` user to write to those directories. For example, if the CRS home directory is `/u01/crs/oracle/product/10.1.0/crs_1`, then enter commands similar to the following:

```
chmod go-w /u01/crs/oracle/product/10.1.0
chmod go-w /u01/crs/oracle/product
chmod go-w /u01/crs/oracle/
chmod go-w /u01/crs/
chmod go-w /u01
```

---



---

#### Notes:

- You must perform this step on all of the nodes that are part of this installation session if you are not installing the CRS software on a cluster file system.
- If you are installing CRS on a Mac OS X-based system, you need to ensure that an admin group user can write to the directories listed in this step. Use commands similar to the following to complete this task:

```
$ sudo chmod go-w /u01/crs/oracle/product/10.1.0
```

Also, the following message might appear when you run the `root.sh` script:

```
You have inserted a disk containing no volumes that MAC OS X
can read. To continue with the disk inserted, click ignore.
```

If you see this message, click Ignore to continue.

---



---

16. In the same terminal window, run the `root.sh` script as the `root` user (or as a user in the admin group on Mac OS X-based systems, with the `sudo` command prefix). Execute the `root.sh` script on *one node at a time*. 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.

After you have run the `root.sh` script on all of the nodes that are part of this installation and when the final `root.sh` script has completed, return to the OUI and click OK on the dialog. The OUI runs the Oracle Notification Server Configuration Assistant and Oracle Private Interconnect Configuration Assistant. These assistants run without your intervention.

---



---

**Note:** If one or more of the configuration assistants fails, you can manually rerun the script created by the OUI to help determine the cause of the failure. You can find the script in the *CRS home/cfgtoollogs* directory. You may need to edit the script to replace placeholders for passwords, to remove steps that completed successfully, and so on.

---



---

17. When the OUI displays the End of Installation page, click Exit to exit the Installer.
18. Unset the `ORACLE_HOME` environment variable.
19. Verify your CRS installation by executing the following `olsnodes` command from the *CRS Home/bin* directory:

```
olsnodes
```

The output from this command should be a listing of the nodes on which CRS was installed, as shown in the following example:

```
$ cd /opt/oracle/app/product/10.1.0/crs_1/bin
$./olsnodes -n

racserver1 1
racserver2 2
```

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 11, "Installing Oracle Database 10g with Real Application Clusters"](#).

---



---

**Note:** The scripts executed during the CRS installation modify the permissions, ownership, and group settings of the OCR and voting disk. After the installation, the values may not be the same as those that you assigned while completing the pre-installation steps for your operating system. Following the CRS installation, the OCR disk should be owned by the `root` user (or by a user in the `admin` group on Mac OS X-based systems), be in the `oinstall` group, and have permissions set to 640; the voting disk should be owned by `oracle`, be in the `oinstall` group, and have permissions set to 644.

---



---

Oracle assumes that when you install Oracle Database 10g Cluster Ready Services software, that you will also install the Oracle Database 10g with RAC software. If there will be a significant delay before you perform phase two of the RAC installation to install Oracle Database 10g with RAC software, then run `CRS_Home/bin/gsdctl start` to start the GSD manually to service the 9.2 SRVCTL tool and assistants. Then before installing the Oracle Database 10g with RAC software, run the command `CRS_Home/bin/gsdctl stop` to stop the GSD.

## 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. Note that this process will only appear on platforms that do not use vendor clusterware with CRS.

---

---

**Note:** the `oproc` process does not run on Linux x86 or Linux Itanium, even though these platforms do not have vendor clusterware.

---

---

- `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 (or by a user in the `admin` group on Mac OS X-based systems) and restarts automatically upon failure.





---

---

## Installing Cluster Ready Services on Windows

This chapter describes the procedures for installing Cluster Ready Services (CRS) on Windows, phase one of the Oracle Database 10g Real Application Clusters installation on Windows-based systems. The topics in this chapter are:

- [Prepare to Install Cluster Ready Services on Windows-Based Systems](#)
- [Using the Oracle Universal Installer to Install Cluster Ready Services on Windows](#)
- [Formatting Drives to Use Oracle Cluster File System after Installing Oracle Database 10g with RAC](#)

### Prepare to Install Cluster Ready Services on Windows-Based Systems

Perform the following procedures prior to initiating the Oracle Universal Installer (OUI) to start phase one of the installation of the Oracle Database 10g with Real Application Clusters (RAC):

- [Verify Cluster Privileges](#)
- [Stop GSD Services from Earlier Releases](#)
- [Storage Configuration Steps for Real Application Clusters](#)

#### Verify Cluster Privileges

Before running the Oracle Universal Installer, from the node where you intend to run the Installer, verify that you have administrative privileges on the other nodes. To do this, enter the following command for each node that is a part of the cluster:

```
net use \\node_name\C$
```

where *node\_name* is the node name.

Oracle Corporation recommends using the same user name and password on each node in a cluster or use a domain user name. If you use a domain user name, log on under a domain with a username and password that has administrative privileges on each node.

#### Stop GSD Services from Earlier Releases

If a GSD from Oracle9i, Release 9.2 or earlier, is running, then stop it before installing Oracle Database 10g CRS by stopping the OracleGSDService in the Services window.

## Storage Configuration Steps for Real Application Clusters

Log in to Windows with Administrative privileges and perform the following steps depending on whether you use the Oracle Cluster File System (OCFS) or raw devices:

- [Using Oracle Cluster File System](#)
- [Using Raw Devices](#)
- [Using Raw Devices for Database When Choosing OCFS for Data Storage During CRS Install](#)

If you plan to use Automatic Storage Management (ASM) for your database files, you only need to perform the actions related to the Oracle home and the Oracle Cluster Registry (OCR) and voting disk storage.

### Using Oracle Cluster File System

If you plan to use OCFS for your Oracle home and datafiles, the following partitions must exist prior to running OUI to install CRS:

- 3 GB, or larger, for the Oracle home
- 3 GB, or larger, for the datafiles

The OCR and voting disk, required by CRS, are also stored in the OCFS datafile directory (*datafile\_disk\cdata\clustername*) where *datafile\_disk* is the OCFS partition and *clustername* is the name of your cluster.

### Using Raw Devices

To use raw devices, create two partitions, one sized at 20 MB for the voting disk and one sized at 100 MB for the OCR. If you are not using OCFS for your datafiles, then you must also create raw partitions for your database files as described in [Appendix C, "Configuring Raw Devices for Real Application Clusters"](#).

### Using Raw Devices for Database When Choosing OCFS for Data Storage During CRS Install

To use raw devices for storing Oracle data files, during the CRS installation on the Select Disk Formatting Options page, choose one of the following options:

- Format one logical drive for software storage
- Do not format any logical drives

These choices create and start the Oracle Object Service on each node, which is required for using raw devices for Oracle database files. If you are using other options that specify data storage on CFS, and if you want to use raw devices, then create and start the Oracle Object Service on each node of the cluster as follows:

1. On each node, run `CRS_home\bin\OracleOBJService.exe /install` to install Oracle Object Service.
2. On each node, start the Oracle Object Service using the Windows Services Control panel.

**See Also:** Microsoft online help for more information about starting services

3. From any of the cluster nodes, run `CRS_home\bin\GUIOracleOBJManager.exe` to create persistent symbolic links to the corresponding raw partitions.

4. After creating the symbolic links, on the GUIOracleObjectManager GUI, perform Options > Commit and Options > Exit

**See Also:** The Oracle9i release 9.2 *RAC Configuration and Setup Guide*, Appendix B for more details

5. From all other nodes, if you already had GUIOracleObjectManager.exe running, then perform Options > Refresh. Otherwise, run *CRS* `home\bin\GUIOracleOBJManager.exe` on these nodes before proceeding to the Oracle Database 10g with RAC installation and before creating your database using DBCA.

---



---

**Note:**

To use raw partitions on Windows 2003, be sure to enable disk automount on all of the nodes of the cluster. This ensures that all of the nodes can access the raw partitions. Do this prior to starting the CRS installation. You can enable automount in one of two ways:

1. Use the diskpart Windows utility and issue the automount enable command from within the diskpart command line.
2. Use the mountvol Windows utility and issue the `mountvol.exe /e` command.

Restart your cluster after performing either of these procedures to ensure that the changes take effect.

---



---

using the instructions you can find in *CRS* `home\relnotes\README_olm.txt`.

## Using the Oracle Universal Installer to Install Cluster Ready Services on Windows

Perform the following procedures to complete phase one, install CRS with the Oracle Universal Installer, of the installation of the Oracle Database 10g with Real Application Clusters (RAC):

---



---

**Note:** During the CRS installation, you may see a message on one or more remote nodes indicating that `lsnodes` cannot find the `oracore10.dll` in the path specified. This is a benign message that you can ignore.

---



---

1. Log in to Windows with Administrative privileges and run the `setup.exe` command on the Oracle Cluster Ready Services Release 1 CD-ROM. This will open the Oracle Universal Installer (OUI) Welcome page.

---



---

**Note:** If you receive a warning to stop all Oracle services during the execution of the OUI, then execute the command

```
Oracle home\bin\localconfig delete
```

where *Oracle home* is the home that is running CSS, to accomplish this.

---



---

2. After you click Next on the Welcome page, the Specify File Locations page provides you with the option to accept the displayed path name for the CRS products or select a different one. You may also accept default directory and path name for the location of your CRS home or browse for an alternate directory and destination. You must select a destination that exists on each cluster node that is part of this installation. Click Next to confirm your choices and proceed to the Language Selection page.
3. Select the language or languages for your CRS installation on the Language Selection page, then click Next for the Cluster Configuration page.
4. The Cluster Configuration page contains predefined node information if the OUI detects that your system has the Oracle9i Release 2 clusterware. Otherwise, the OUI displays the Cluster Configuration page without predefined node information.

Provide your own cluster name if you do not wish to use the name provided by the OUI. Note that the selected cluster name must be globally unique throughout the enterprise and its allowable character set is the same as that for hostnames, that is, underscores (`_`), hyphens (`-`), and single-byte alphanumeric characters (a to z, A to Z, and 0 to 9).

---

---

**Note:** If your cluster has Oracle9i Release 2 clusterware installed, then you must enter the same cluster name that you used for the Oracle9i Release 2 installation or else the installation will fail.

---

---

Enter a public and a private host name for each node. Neither the public nor the private name should have a domain qualifier. When you enter the public host name, use the primary host name of each node, that is, the name displayed by the `hostname` command. The private node refers to an address that is only accessible by the other nodes in this cluster, and which Oracle uses for Cache Fusion processing. You may enter either a private host name or a private IP address for each node.

Click Next after you have entered the cluster configuration information. This saves your entries and opens the Specify Network Interface Usage page.

5. In the Specify Network Interface Usage 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`. Click Next when you have made your selections to open the Select Disk Formatting Options page.
6. On the Select Disk Formatting Options page, you indicate what OCFS you want the OUI to format for you.

---

---

**Note:** To use ASM for your database storage, then for this step select the Do not format any logical drives option.

---

---

The formatting options are as follows:

- Format two logical drives for data and software storage. In this case, the OUI creates two cluster file systems, one for the database files and one for the Oracle home. The OCR and voting disk are created in the datafile directory.

- Format one logical drive for software storage. In this case, the OUI creates one cluster file system for the Oracle home and requires two additional partitions, one for the OCR and one for the voting disk. If you select this option, then you must have configured the partitions as described in "[Creating Partitions and Logical Drives](#)" on page 8-18 before proceeding with the installation. These partitions will be stamped with "ocrcfg" for the OCR and "votedsk" for the voting disk.
- Format one logical drive for data file storage. In this case, the OUI creates one cluster file system for the database files which is also used to store the OCR and voting disk.
- Do not format any logical drives. In this case, the OUI requires two partitions, one for the OCR and one for the voting disk. If you select this option, then you must have configured the partitions as described in "[Creating Partitions and Logical Drives](#)" on page 8-18 before proceeding with the installation. These partitions will be stamped with "ocrcfg" for the OCR and "votedsk" for the voting disk.

Click Next after making your selection on the Select Disk Formatting Options page. The next page that the OUI displays depends on your selection on the Select Disk Formatting Options page, as described in the following step.

---

---

**Note:** If you are installing on a cluster with an existing cluster file system from an earlier release of Oracle, then the OCR and voting disk will be stored in that file system. In this case, you do not require new partitions for the OCR and voting disk, even if you do not format a logical drive for data file storage.

---

---

7. If you selected the "Format two logical drives for data and software storage" option on the Select Disk Formatting Options page, then you must complete the Select Software Storage Drive page and the Select Data Storage Drive page as described in Steps a and b respectively.

If you selected the "Format one logical drive for software storage" option on the Select Disk Formatting Options page, then you must complete the Select Software Storage Drive page, the Disk Configuration - Oracle Cluster Registry (OCR) page, and the Disk Configuration - Voting Disk page as described in Steps a, c, and d respectively.

If you selected the "Format one logical drive for data file storage" option on the Select Disk Formatting Options page, then you must complete the Select Data Storage Drive page as described in Step b.

If you selected the "Do not format any logical drives" option on the Select Disk Formatting Options page, then you must complete the Disk Configuration - Oracle Cluster Registry (OCR) page and the Disk Configuration - Voting Disk page as described in Steps c and d respectively.

- a. On the Select Software Storage Drive page, choose a shared drive to see a list of available partitions on that drive. Choose a partition with sufficient space to hold your Oracle home and select the partition's disk number and partition number from the list. Click Next to proceed.
- b. On the Select Data Storage Drive page, choose a shared drive to see a list of available partitions on that drive. Choose a partition with sufficient space to hold your database files and select the partition's disk number and partition number from the list. Click Next to proceed.

- c. On the Disk Configuration - Oracle Cluster Registry (OCR) page, choose a partition with sufficient space to hold your OCR and select the partition's disk number and partition number from the list. Click Next to proceed.
- d. On the Disk Configuration - Voting Disk page, choose a partition with sufficient space to hold your voting disk and select the partition's disk number and partition number from the list. Click Next to proceed.

---

---

**Note:** The OUI pages described in this step display logical drives from which you must make your selections. To be valid for selection, a logical drive must be located on a disk without a primary partition.

---

---

8. After you click Next, the OUI checks whether the remote inventories are set. If they are not set, then the OUI sets up the remote inventories by setting registry keys. The OUI also verifies the permissions to enable writing to the inventory directories on the remote nodes. After completing these actions, the OUI displays a Summary page that shows the cluster node information along with the space requirements and availability. Verify the installation that the OUI is about to perform and click Finish.
9. When you click Finish, the OUI installs the OCFS and CRS software on the local node and validates the installation again. The OUI will also create any required OCFS file systems. After validating the installation, the OUI completes the CRS software installation and configuration on the remote nodes, including the execution of required configuration assistants.

---

---

**Note:** If one or more of the configuration assistants fails, you can manually rerun the script created by the OUI to help determine the cause of the failure. You can find the script in the *CRS* *home*\*cfgtoollogs* directory. You may need to edit the script to replace placeholders for passwords, to remove steps that completed successfully, and so on.

---

---

10. For installations of Cluster Ready Services on a system that also contains Oracle9i Real Application Clusters, note these additional considerations and complete the steps as necessary:
  - Restart all of the newly installed Oracle Database 10g release 1 (10.1) cluster member nodes. You can restart one node at a time so that availability of the Oracle9i databases is not disrupted.
  - If there will be a significant delay before you perform phase two of the RAC installation to install Oracle Database 10g with RAC software, then execute the command `CRS_home\bin\gsdctl start` to start the GSD manually to service the 9.2 SRVCTL tool and assistants. Then, before configuring the Oracle Database 10g with RAC software, run the command `CRS_home\bin\gsdctl stop` to stop the GSD.

At this point, you have completed phase one, the CRS installation and are ready to install the Oracle Database 10g with RAC as described in [Chapter 11, "Installing Oracle Database 10g with Real Application Clusters"](#).

## Formatting Drives to Use Oracle Cluster File System after Installing Oracle Database 10g with RAC

If you install the Oracle Database 10g with RAC, and later want to install OCFS, then execute the `ocfsformat.exe` command from the `crs_home\cfs` directory using the following syntax:

```
crs_home\cfs\ocfsformat.exe /l drive_letter /c clustersize [/v volume_label] [/f]
```

where:

- *drive\_letter* is the drive on which you want to format the OCFS
- *clustersize* is the size of the partition in kilobytes
- *volume\_label* is an optional volume label





---

---

# Installing Oracle Database 10g with Real Application Clusters

This chapter describes phase two of the 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 procedures in this chapter apply to both UNIX- and Windows-based systems. This chapter contains the following topics:

- [Selecting a Database Configuration Type](#)
- [Installation Setup Procedures](#)
- [Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer](#)
- [Installation on Windows-Based Systems with the Minimum Memory Requirements](#)
- [De-Installing Real Application Clusters Software](#)

## 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 Advanced 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 Advanced configuration, then you can use the Database Configuration Assistant (DBCA) to create the database as described in [Chapter 12](#). Oracle recommends that you use the DBCA to create your database.

You can also select the Advanced 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 Advanced template to create a database.

Oracle recommends that you use one of the preconfigured database options or use the Advanced option and the DBCA to create your database. However, if you want to configure your environment and create your database manually, then select the "Do not create a starter database" configuration option and refer to the manual database creation procedures posted at <http://www.oracle.com/technology>.

## Configuration Type Descriptions

The configuration type that you select, as described in [Table 11-1](#), determines how you proceed.

**Table 11-1 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 configures your RAC database. | Minimal input required. You can create your database more quickly than with the Advanced type.        |
| Advanced                                                    | 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. |
| Do not create a starter database                            | 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 Select Database Configuration page, then complete one of the platform-specific procedures in this chapter under the heading "[Installation Setup Procedures](#)" on page 11-3. These three configuration types use preconfigured templates. After you complete these procedures, the Oracle Network Configuration Assistant (NETCA) 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 choose raw devices on the Specify Database File Storage Option page, then the DBCA verifies that you configured the raw devices for each tablespace.

If you select the Advanced configuration, then you must enter specific information as described in the next section.

### Using the Advanced Configuration Type

When you select the Advanced configuration type, the OUI runs the DBCA which displays four preconfigured database template choices:

- General Purpose
- Transaction Processing
- Data Warehouse
- Advanced

The first three templates create a database that is optimized for that environment. You also can customize these templates. The Advanced type, however, creates a database without using preconfigured options.

The following section provides more detail about OUI and DBCA processing when creating a RAC database.

## Behavior of the OUI, the DBCA, and Other Assistants During Installation

After installation, the OUI starts the NETCA. After the NETCA 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

Complete one of the following platform-specific installation setup procedures:

- [Installation Setup Procedures for UNIX-Based Systems](#)
- [Installation Setup Procedures for Windows-Based Systems](#)

### Installation Setup Procedures for UNIX-Based Systems

If the user who installed Cluster Ready Services (CRS) is different from the user who installs RAC, then all of the user-level pre-install steps must be completed.

---

**Note:** On Mac OS X systems, you must run the Installer from the terminal session where you logged in as the oracle user and set the user's environment (described in [Chapter 6, "Pre-installation Tasks for RAC on Mac OS X"](#)).

---

1. Verify user equivalence by running 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 run 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.

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 CRS installation.

---

You can use your Oracle9i database language and territory definition files with the Oracle Database 10g release 1 (10.1) that you are about to install. To enable this functionality, you must run the OUI from a command line, described in Step1 in ["Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer"](#) on page 11-5, and use following statement to set the `b_cr9idata` variable to `true`:

```
runInstaller oracle.rsfnlsrtl_rsf:b_cr9idata=true
```

## Installation Setup Procedures for Windows-Based Systems

Perform the following preliminary steps for your RAC installation:

1. Perform a network connectivity test to ensure that all the 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.
2. Log on to one of the nodes in your cluster as a member of the Administrators group. Log on as a member of the Domain Administrators group if you are installing on a Primary Domain Controller (PDC) or a Backup Domain Controller (BDC).
3. Delete the `ORACLE_HOME` environment variable if it exists. Refer to your Microsoft online help for more information about deleting environment variables.

---

**Note:** The `ORACLE_HOME` is automatically set in the registry. Manually setting this environment variable prevents installation.

---

4. If you are installing in an existing Oracle Database 10g release 1 (10.1) home, then stop all Oracle services that are running from that existing home. These services will have names with names such as `OraclehomeName` (this includes the Oracle listener, `OraclehomeNameTNSListener`), `OracleServiceSID`, or `OracleASMServicesSID`. Be careful not to stop services running from other homes, particularly `OracleCSService`, `OracleEVMService`, and `OracleCRService`.

Refer to your Microsoft online help for more information about stopping services.

You can use your Oracle9i database language and territory definition files with the Oracle Database 10g release 1 (10.1) that you are about to install. To enable this functionality, you must run the OUI from a command line, described in Step1 in

"Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer" on page 11-5 and use following statement to set the `b_cr9idata` variable to `true`:

```
setup.exe oracle.rsfnlsrtl_rsf:b_cr9idata=true
```

## Installation of Oracle Database 10g with RAC Using the Oracle Universal Installer

Perform the following procedures to install the Oracle Database 10g software with RAC.

1. On UNIX-based systems, run the `runInstaller` command from the top-level directory of the Oracle Database 10g Release 1 CD-ROM.

On Windows-based systems, insert the Oracle Database 10g Release 1 CD-ROM and the Autorun screen should appear. If the Autorun screen does not appear, then:

- a. Click Start > Run
- b. Enter the following command:

```
DRIVE_LETTER:\autorun\autorun.exe
```

If you are installing the software from a hard drive, then navigate to the parent directory of the directory where the Oracle Database 10g product is stored and run the `setup.exe` program.

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 `products.xml`. Enter the Oracle home name and location in the target destination and click Next.

---

---

**Note:** The Oracle home name and path 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 in that directory. If the Oracle home exists, if the Oracle home is registered with the OUI inventory, and if the Oracle home was created with a cluster install, 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 Select Installation Type page described in Step 6.

---

---

**Note:** Even if you are using a clustered file system, then the Selected Nodes page should display all the nodes you intend to use for your RAC databases.

---

---

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* or to the *Oracle Database Platform Guide for Windows* to perform a single-node, non-RAC installation on a UNIX or Windows cluster, or 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 installation log file. On UNIX-based systems, the installation log file is:

`OraInventory/logs/installActionsdate_time.log`

On Windows-based systems, the installation log file is:

`system_drive:\Program Files\Oracle\Inventory\logs\installActionsdate_time.log`

You can either make corrections and click Next, or you can deselect nodes that have errors and click Next. When you click Next, the OUI displays the Select Installation 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 Select Installation Type page you can select Enterprise Edition, Standard Edition, or Custom Install type. If you select Enterprise Edition or Standard Edition, then your installation will include the components related to the chosen edition. If you select the Custom install, then select the individual components that you want to install from the list of available components. Note that if you have purchased a Standard Edition license and you perform a Custom installation, then only install the products covered by the Standard Edition license.

After you make your selection and click Next on the Select Installation Type page, the OUI displays one of the following pages, depending on your configuration and selections:

- Product-specific Prerequisite Checks page: This page is displayed during installations on UNIX-based systems and is described in Step 7.
  - Upgrading an Existing Database page: This page is displayed if the OUI detects a previous release of the database and is described in Step 8.
  - Select Database Configuration page: This page is displayed if no previous release of the database exists and is described in Step 10.
7. The Product-specific Prerequisite Checks page verifies the operating system kernel parameters or attributes and calculates the `ORACLE_BASE` location. Click Next and the next page that the OUI displays depends on your configuration and selection options as follows:
    - Upgrading an Existing Database page: This page is displayed if the OUI detects a previous release of the database and is described in Step 8.
    - Available Products Components page: This page is displayed if you selected a Custom install and is described in Step 9.
    - Select Database Configuration page: This page is displayed if you did not select a Custom install and there is no database to upgrade. This page is described in Step 10.
  8. On the Upgrading an Existing Database page, you can choose to upgrade one or more existing databases or continue with the installation without performing any upgrades. If you wish to upgrade an existing database, then select the check box labelled "Upgrade an Existing Database" and then select the databases to upgrade from the list provided. If you prefer not to upgrade any of the listed databases, then leave the "Upgrade an Existing Database" check box unchecked.

---

**Note:** When DBUA displays a list of databases to upgrade, it may display duplicate entries for the same database. This occurs on Windows systems for Oracle9i release 9.2 databases created using DBCA that have database names containing non-uppercase characters. In this case, an additional entry will be displayed consisting of the database name that contains upper-case characters. You can choose either of the entries to proceed with upgrading the database.

---

If you selected "Upgrade an Existing Database", then the OUI continues with the Summary page, described in Step 18, when you click Next. If you did not select "Upgrade an Existing Database", then the OUI continues with the Select Database Configuration page, described in Step 10, when you click Next.

9. On the Available Products Components page, select the components you wish to install and click Next. This opens the Component Locations page in some cases,

such as when insufficient disk space is available. Choose the location for your install and click Next. The OUI will display the Privileged Operating System Groups page described in Step 17.

10. The Select Database Configuration page, provides options for you to create a preconfigured database as part of the installation process or to install the software without creating a database. If you want to create a database as part of your installation, then select General Purpose, Transaction Processing, Data Warehouse, or Advanced. If you only want to install the software, then select "Do not create a starter database". Click Help for more information about these choices.

---

---

**Note:** If you are installing on a Linux system which uses an ASM library driver, then you must select either the Advanced option or a software only install. Additionally, you must specify `ORCL:*` as the disk discovery path when asked for this information in the DBCA.

---

---

If you choose to create a General Purpose, Transaction Processing, or a Data Warehouse database, then when you click Next, the OUI displays the Specify Database Configuration Options page, described in Step 11. If you choose the Advanced option or select "Do not create a starter database", then when you click Next, the OUI displays the Privileged Operating System Groups page, described in Step 17.

---

---

**Note:** If you perform an Advanced or a software only installation on a Windows cluster, then the OUI does not perform Steps 19 through 23 and, therefore, does not run the Virtual IP Configuration Assistant (VIPCA), the NETCA, and the DBCA. To run these configuration tools manually on Windows, complete the additional processing described in Step 25.

---

---

11. On the Specify 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. Note that a SID prefix cannot exceed 5 characters on UNIX platforms and 61 characters on Windows platforms. Also select a database character set and, on UNIX-based systems only, select any database examples to install from the Sample Schemas.

When you click Next on the Specify Database Configuration Options page, the OUI displays the Select Database Management Option page.

12. On the Select Database Management Option page, if you are installing on a UNIX-based system and you have already completed the Grid Control Management Agent installation, then you can select either Grid or Local Database control. Otherwise, only Local Database control for database management is supported for RAC. When you use the local Database Control, you can choose the email option and enter the outgoing SMTP server name and email address.



---



---

**Note:** The option to configure central management for Enterprise Manager is not available during RAC installation on Windows. Also not supported on Windows is the use of standalone Enterprise Manager Configuration Assistant or Database Configuration Assistant to configure central management for RAC. If you want central management for the installed RAC database, then you will have to discover the RAC database target manually from Grid Control after the installation.

---



---

If you perform an installation that does not include Enterprise Manager, for example, a custom software install without Enterprise Manager, an installation with no Enterprise Manager configuration, or a database creation with your own scripts, then you can configure Enterprise Manager later with the OUI, the DBCA, or the Enterprise Manager Configuration Assistant (EMCA) utility.

**See Also:** *Enterprise Manager Grid Control Installation and Basic Configuration* for details about installing Grid Control with the OUI utility and *Enterprise Manager Advanced Configuration Guide* for details about installing Database Control with the DBCA and the EMCA utilities.

When you click Next on the Select Database Management Option page, the OUI displays the Specify Database File Storage Option page.

13. Select a data storage option on the Specify Database File Storage Option page.

---



---

**Note:** If you are installing RAC from the Standard Edition, then you must use ASM for your database storage. Even if you can click other database storage options on this page, they are only supported by the Enterprise Edition and you should not select them.

---



---

**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 as follows:

- On UNIX-based systems, the default partitions dependent on the operating system, as follows:

| Operating System | Default Search String |
|------------------|-----------------------|
| AIX              | /dev/rhdisk*          |
| Mac OS X         | /dev/rdisk*           |
| HP-UX            | /dev/rdisk/*          |
| Tru64 UNIX       | /dev/rdisk/*          |
| Linux            | /dev/raw/*            |
| Solaris          | /dev/rdisk/*          |

- On Windows-based systems, the default partitions for ASM are:

\\.\orcldisk\*

---

---

**Note:** The only partitions that the OUI displays for Windows systems are logical drives that are on drives that do not contain a primary partition and have been stamped with `asmt001`.

---

---

**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 Specify Database File Storage Option page, the OUI displays the Specify Backup and Recovery Options page.

14. If you enable backup on the Specify Backup and Recovery Options page, then you can choose File System or ASM. You will also enter the user name and password. When you click Next on the Specify Backup and Recovery Options page, the next page the OUI displays depends on your configuration and selection options:
  - If you have selected ASM as a storage option, then the next page is either the Select ASM Disk Group page, if there is an existing ASM instance and disk group available, or the Configure Automatic Storage Management page, described in Step 15, if you need to define your ASM environment.
  - If you have not selected ASM as a storage option, then the next page is the Specify Database Schema Passwords page, described in Step 16.
15. The Configure Automatic Storage Management page lists the available disk partition locations. Select the disks that you prepared for your ASM disk groups. To continue to the Specify Database Schema Passwords page, click Next.
16. On the Specify Database Schema Passwords 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 Specify Database Schema Passwords page, the next page that the OUI displays depends on your configuration and selection options as follows:
  - On UNIX-based systems, if the user is not a member of the `dba` group, then the next page is the Privileged Operating System Groups page, described in Step 17.
  - On UNIX-based systems, if the user is a member of the `dba` group, then the next page is the Summary page, described in Step 18.
  - On Windows-based systems, the next page is the Summary page, described in Step 18.
17. The Privileged Operating System Groups page, is for UNIX systems only. On this page, you enter the group name for the `SYSDBA` and `SYSOPER` users. When you click Next on the Privileged Operating System Groups page, the OUI displays the Create Database page, which you can ignore if you are upgrading. Then, if you are installing on a Tru64 UNIX system, the OUI displays the Choose JDK Home Directory page. Finally, the OUI proceeds to the Summary page.
18. 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.

---

---

**Note:** On Tru64 UNIX systems, an additional page, the Choose JDK Home Directory page will appear before the Summary page. Provide the required directory information and click Next to proceed to the Summary page.

---

---

During the installation, unless you are placing your Oracle home on a clustered file system, the OUI copies software to the local node and then copies the software to the remote nodes. On UNIX-based systems, the OUI then prompts you to run the `root . sh` script on all the selected nodes to start the VIPCA and display the VIPCA Welcome page. Before you run `root . sh` on UNIX-based systems other than Mac OS X, make sure that your display environment variables are properly set. On Windows-based systems, the VIPCA runs from OUI as a configuration assistant unless you performed an Advanced or software only installation. In the case of a Windows installation without a preconfigured database, Steps 19 through 23 will not be performed and you should proceed with the manual processing described in Steps 24 and 25.

19. Review the information on the VIPCA Welcome page, click Next, and the VIPCA displays the Public Network Interfaces page.

---

---

**Note:** Depending on the size of your cluster as well as other factors, when running the Virtual Internet Protocol Assistant (VIPCA) in a shared Oracle home on Oracle Cluster File System (OCFS) in a Windows environment, the GSD may be slow to start. A time out may occur after 15 minutes or more with the error Failure in starting GSD resource on *node*. Workaround: If the error and time out occurs, then select Ignore to continue. VIPCA will proceed much faster and complete its operations more quickly.

---

---

20. 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.
21. 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.
22. On UNIX-based systems, run the `root . sh` procedure, one node at a time, on each of the nodes that is part of this installation. The VIPCA will not run again on the remote node because the node applications for the remote node are already configured.
23. 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:
  - NETCA
  - DBCA

At the end of the database creation, DBCA configures Database Control and starts the services. Database Control is immediately available to administer and monitor your newly installed Oracle Database 10g with Real Application Clusters environment.

24. On Windows-based systems only, ensure that the path name for your new Oracle home is defined across the cluster by completing the following procedure on each remote node that is part of your current cluster installation:
  - a. On each remote node, navigate to Start > Control Panel > System > Advanced > Environment Variables
  - b. In the "System variables" dialog, select the Path variable and ensure that the value for the Path variable contains *oracle\_home\bin*, where *oracle\_home* is your new Oracle home. If the variable does not contain this value, then click Edit and add this value to the start of the path variable definition in the Edit System Variable dialog and click OK.
  - c. Click OK in the Environment Variables page, then click OK in the System Properties page, and then close the Control Panel
25. On Windows-based systems only, if you selected the "Do not create a starter database" option in Step 10, and you decide to run NETCA and or DBCA after the installation, then you must first configure and start the VIPs, GSD and ONS applications by running VIPCA. To start VIPCA, run *vipca.bat* from the *%ORACLE\_HOME%\bin* directory. You run VIPCA only once for the cluster.

You have completed the second and final phase of the installation. Proceed to [Chapter 13, "Real Application Clusters Post-Installation Procedures"](#) to perform the post-installation tasks.

---

**Notes:**

- If you need to change the VIP on a RAC node, then you should use the command

```
srvctl modify nodeapps -A new_address
```

where *new\_address* is defined as explained in Appendix B of the *Oracle Real Application Clusters Administrator's Guide*.

- Use the DBCA and the OUI to de-install RAC.
- 

## Installation on Windows-Based Systems with the Minimum Memory Requirements

Installations of RAC on nodes in Windows-based systems with 256 MB of RAM and 500 MB of virtual memory have the following limitations:

- Computers with 256 MB of memory are not able to run Oracle Database Upgrade Assistant (DBUA), DBCA, or NETCA during an Oracle Universal Installer installation session.
- Depending on how many applications are running on the computer, you may need to further increase the paging file size or reduce the size of the System Global Area (SGA) if you run out of virtual memory. Note that if temporary files and the paging file are both stored on the same physical drive, then a situation can occur where the space requirements for one can limit the size of another. If your system has limited free space, then first install the Oracle Database software. After the

installation is finished, run NETCA to configure listeners and then DBCA to create a database.

On computer systems that barely meet the minimum memory and virtual memory requirements, 256 MB and 500 MB respectively, do not install the database. Follow these guidelines:

- Select Enterprise Edition Installation and deselect "Do not create a starter database"
- Select Custom Installation, select "Do not create a starter database" from the Select Database Configuration page.
- Select Advanced Installation, select the Custom installation type from the Select Installation Type page, and select No on the Create Database page when prompted to create the database.
- Cancel DBCA from the Configuration Assistants page.

After installation, run the appropriate configuration assistant for your needs:

- To create a new database, run DBCA from the Start Menu. Choose Start > Programs > Oracle - *HOME\_NAME* > Configuration and Migration Tools > Database Configuration Assistant.
- To upgrade an existing database, run DBUA from the Start Menu. Choose Start > Programs > Oracle - *HOME\_NAME* > Configuration and Migration Tools > Database Upgrade Configuration Assistant.

To ensure that all nodes in your cluster use the new Oracle home, ensure that the path name for the new Oracle home is defined on each node in the cluster by completing the following procedure:

1. On each node, navigate to Start > Control Panel > System > Advanced > Environment Variables
2. In the "System variables" dialog, select the Path variable and ensure that the value for the Path variable contains *oracle\_home*\bin, where *oracle\_home* is your new Oracle home. If the variable does not contain this value, then click Edit and add this value to the start of the path variable definition in the Edit System Variable dialog and click OK.
3. Click OK in the Environment Variables page, then click OK in the System Properties page, and then close the Control Panel.

## De-Installing Real Application Clusters Software

A successful de-installation of Oracle software from a cluster requires you to run the OUI from the same node that was used for the initial installation. This is required for the RAC and for the CRS software de-installations. To determine the node used for the installation, perform the following steps:

1. Navigate to the Oracle home directory if you are de-installing RAC software or to CRS home directory if you are de-installing CRS software.
2. If you are on a UNIX-based system, then execute the command `./runInstaller`. If you are on a Windows-based system, then run the `setup.exe` program.
3. Click "Installed Products..."
4. Expand the node for the Oracle home or the CRS home you wish to de-install.

5. In the cluster nodes subtree that lists the nodes on which this home is installed, the first node listed is the node on which the install was performed. This is the node where you should run OUI in the following procedures.

Once you have identified the node where your software was installed, perform the following procedures to de-install Oracle Database 10g RAC and CRS software. On both UNIX- and Windows-based systems, you must de-install the Oracle database software first before de-installing the Cluster Ready Services (CRS) software. These steps to complete these de-installation tasks are described in the following sections:

- [De-Installing Oracle Database 10g RAC Software](#)
- [De-Installing Cluster Ready Services](#)

---

---

**Note:** These sections describe a complete de-installation of the RAC, ASM, and CRS software where RAC and ASM share the Oracle home and no other Oracle home exists.

If you have multiple Oracle homes on the cluster, then check for any dependencies that might affect your other databases. Such dependencies can include listeners, ASM instances, and so on that run in the Oracle home to be deleted. To identify dependencies, on UNIX-based platforms, review the `oratab` file to identify common Oracle homes and on Windows-based platforms, refer to the contents of your Oracle Cluster Registry (OCR).

---

---

**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 and for information about viewing OCR content

## De-Installing Oracle Database 10g RAC Software

This section describes the procedure to de-install the Oracle Database 10g RAC software. Before you perform these steps, consider making a backup of any databases that run from the Oracle home you are about to delete. You should then stop any instances and processes on all nodes, including services on Windows-based systems, that depend on the software being de-installed.

1. Delete all of the databases that are dependent on the Oracle home that you are deleting by using the DBCA "Delete a database" option.
2. If ASM runs from the same Oracle home that you are de-installing, then ensure that there are no other database dependencies on this ASM instance and remove the ASM configuration by logging on as the `oracle` user and completing the following steps:
  - a. Connect to the ASM instance and run the following command on all of the nodes that are affected by the de-installation of this Oracle home to identify the database instances that are using ASM:

```
SQL> select INSTANCE_NAME from V$ASM_CLIENT;
```

---



---

**Note:** This command only lists the database instances that are running. It is possible that other instances are associated with the ASM instance, but they are not currently running. If you removed a database from this Oracle home but the output from the command shows that this ASM instance is supporting a database instance in another Oracle home, then do not remove the ASM instance or the Oracle home.

---



---

- b. For each instance listed in the output of the statement that you ran in Step a, stop the respective databases.
- c. Oracle recommends that you backup the database files for all of the databases that are currently using this ASM instance.
- d. Using your connection to the ASM instance, run the following command to identify the disk groups that are managed by the ASM instance:

```
SQL> select * from V$ASM_DISKGROUP;
```

- e. For each diskgroup listed in the output of the statement that you ran in Step d, run the following command to drop the disk groups:

```
SQL> drop diskgroup diskgroup_name including contents;
```

where *diskgroup\_name* is the name of the diskgroup.

- f. Shutdown ASM on all of the nodes on which this Oracle home exists. Do this by running the command `srvctl stop asm -n node_name` for all nodes on which this Oracle home exists.
- g. To remove the ASM configuration, run the following command for all nodes on which this Oracle home exists:

```
srvctl remove asm -n node_name
```

where *node\_name* is the name of a node from which you want to remove the ASM instance.

- h. If you are de-installing from a Linux-based system with ASMLIB, then run the following procedure. First, run this command to list the disks that you need to delete:

```
oracleasm listdisks
```

Second, run the following command to delete every disk listed by the previous command:

```
oracleasm deletedisks
```

Third, re-run the `oracleasm listdisks` command to confirm that the disks are all deleted. Repeat this command on each of the RAC cluster nodes confirm that the disks have been deleted from all the nodes.

Fourth, unless you are de-installing the Oracle software from a Mac OS X-based or Windows-based system, run the following commands on all nodes of your RAC cluster as the `root` user:

```
/etc/init.d/oracleasm stop
/etc/init.d/oracleasm disable
```

If you are de-installing the Oracle software from a Mac OS X-based system, run the following commands as a user in the `admin` group:

```
sudo /etc/init.d/oracleasm stop
sudo /etc/init.d/oracleasm disable
```

If you are de-installing the Oracle software from a Windows-based system, then run the `asmtoolg` wizard from the `crs_home/BIN` directory and use the GUI to delete the stamps on the ASM disk partitions. Select the Delete labels option and click Next. Select the disks from which you want to remove the ASM stamps, click Next, and click Finish to complete the operation.

- i. If you are using a shared cluster file system for your Oracle home, then on UNIX-based systems run the following commands on the local node:

```
rm -f $ORACLE_HOME/dbs/*ASM*
rm -r $ORACLE_BASE/admin/+ASM
```

On Windows-based systems run the following commands on the local node:

```
delete %ORACLE_HOME%\database*ASM*
delete %ORACLE_BASE%\admin\+ASM
```

You may need to remove subordinate files or subdirectories before these commands can complete successfully.

- j. If you are not using a shared cluster file system for your Oracle home, then run the commands from the previous step, Step i, on *each* node on which the Oracle home exists.
- k. On UNIX-based systems, remove the `oratab` entries that begin with the characters `+ASM` for each ASM instance on which you performed the previous step, step j. On Windows-based systems, run the following command on each node that has an ASM instance:

```
oradim -delete -asmsid +ASMnode_number
```

where `node_number` is the node identifier.

3. If the listener runs from this Oracle home, then use the NETCA to remove the listener and its configuration.
4. If the CRS node applications for VIP, ONS, and GSD were created in this Oracle home, then you can remove them and re-create them in another Oracle Database 10g Oracle home, or you can modify the Oracle home for these applications so that they use an alternate Oracle Database 10g Oracle home. Both of these alternative procedures are described in the following two points.

Perform *one* of the following procedures:

- a. Stop and remove the CRS node applications on each node that is associated with the Oracle home that you are deleting. Do this by running the command `svrctl stop nodeapps -n node_name` for all of the nodes that are affected by the deletion of the Oracle home. Then remove the CRS node applications by running one of the following platform-specific commands:

- On UNIX-based systems, run the command:

```
$ORACLE_HOME/install/rootdeletenode.sh
```

- On Windows-based systems, run the command:

```
svrctl remove nodeapps -n node_name
```



where *node\_name* is the node name, for each node in the cluster.

Respond to any operating-system prompts to confirm your operations for each node. If other Oracle Database 10g RAC Oracle homes exist, then re-create the node applications in that Oracle home by running the command `srvctl create nodeapps`.

- b. Alternatively, you can modify the CRS node application's Oracle home by running the following command:

```
srvctl modify nodeapps -o 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 and for details of the SRVCTL command syntax.

5. After you complete the previous steps in this procedure, run the OUI and on the Welcome page, click Deinstall Products to display the list of installed products on which you can select the Oracle home to de-install.

---



---

**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 de-install RAC with the OUI and want to perform another RAC installation, then you must start a new OUI session.

---



---

## De-Installing Cluster Ready Services

De-install each Oracle Database 10g RAC home by running the procedure in the previous section, "[De-Installing Oracle Database 10g RAC Software](#)". Then complete the de-installation by removing the CRS software using one of the following operating system-specific procedures:

- [De-Installing Cluster Ready Services from UNIX Environments](#)
- [De-Installing Cluster Ready Services from Windows Environments](#)

### De-Installing Cluster Ready Services from UNIX Environments

Perform the following steps to de-install the CRS software from a UNIX environment:

1. Run the command `CRSHome/install/rootdelete.sh` to disable the CRS applications that are 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` as the first argument, otherwise use `local` as the first argument. For the second 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 on each node of the cluster from which you want to de-install CRS.
2. Run the script `CRS Home/install/rootdeinstall.sh` on a local node to remove the OCR.
3. 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 de-install.

---

---

**Note:** A node is a local node if you plan to run Steps 2 and 3 on that node.

---

---

### De-Installing Cluster Ready Services from Windows Environments

This section describes the following de-installation procedures for removing CRS from Windows environments:

- [De-Installing Cluster Ready Services from Windows Environments with No Previous Cluster Software Versions](#)
- [De-Installing Oracle Cluster Ready Services from Windows Environments, with Clusterware Downgrade to 9.2](#)

**De-Installing Cluster Ready Services from Windows Environments with No Previous Cluster Software Versions** Perform the following steps to de-install Oracle 10g CRS software from a Windows environment:

1. 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 de-install.
2. If you have services with names such as `OracleCRSTokenname`, then remove them by running the following command:  

```
crsuser remove user_name
```

where `user_name` is a user name.
3. Shut down and restart each node that is part of your cluster.
4. If you are not using a cluster file system, then on *each* node, use Windows Explorer to delete the Oracle directory, its subdirectories, and their contents.

**De-Installing Oracle Cluster Ready Services from Windows Environments, with Clusterware Downgrade to 9.2** Perform the following procedures to de-install the Oracle 10g CRS software from a Windows environment that also has 9.2 RAC:

1. Run `CRS_Home\bin\GuiOracleOBJManager.exe` to make sure that the symbolic link named `srvcfg` exists and points to a disk partition, if you are not using OCFS to store the OCR.
2. If Oracle9i release 9.2 clusterware uses Oracle CFS (Cluster File System), then make sure that on all nodes the registry key `HKLM\SOFTWARE\Oracle\osd9i\ocr` exists, and has a string value `CfsOcrRoot` that points to a release 9.2 OCR (CFS) location.
3. Run `CRS_Home\bin\ocrconfig -downgrade` to downgrade the Cluster Registry to a release 9.2 OCR.
4. On each node, copy `CRS_Home\cfs\OcfsFindVol.exe` to `%SYSTEMROOT%\system32\osd9i\cfs`.
5. Run `CRS_Home\oui\bin\setup.exe` to invoke Oracle Universal Installer (OUI) and on the Welcome page click **Deinstall products** to list all the installed products. Select **Oracle CRS home name** from the displayed products and click **Remove** to deconfigure and de-install the product.
6. On each node, run `%SYSTEMROOT%\system32\osd9i\olm\OracleOBJService.exe /install` to re-install the Oracle 9.2 object service. Then start Oracle Object Service.

7. On each node, run `%SYSTEMROOT%\system32\osd9i\cfs\OcfsFindVol.exe /i:%SYSTEMROOT%\system32\osd9i\cfs\OcfsFindVol.exe` to re-install the Oracle Cluster Volume service. Then start the `OracleClusterVolumeService`.
8. From a 9.2 RAC Oracle Home on each node run the command `ORACLE_HOME\bin\gsdservice.exe -install`. Then start the `OracleGSDService`.
9. On each node, copy `%SYSTEMROOT%\system32\osd9i\orafencedrv.sys` `%SYSTEMROOT%\system32\drivers\orafenceservice.sys`



---

---

# 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
- Configure Enterprise Manager Grid Control
- 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 12-2 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.

On Windows-based systems, if you plan to use ASM storage, then you must have created logical partitions without primary partitions on the same drive and deleted the

drive letters for these partitions on all nodes or stamped these partitions with `asmtool`.

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, connect to one of your nodes with Oracle RAC installed and then:

- On UNIX-based platforms, enter the `dbca` command from the `$ORACLE_HOME/bin` directory
- On Windows-based platforms, choose Start > Programs > Oracle - [HOME\_NAME] > Configuration and Migration Tools > Database Configuration Assistant

## 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 `oraInventory` is not corrupted and, on UNIX-based systems, that the OUI inventory, `oraInst.loc`, is located in the correct directory. On AIX, Linux x86, and Linux Itanium, `oraInst.loc` should be found in `/etc` and on other UNIX-based systems `oraInst.loc` should be in `/var/opt/oracle`. 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 begin with an alphabetical character and contain no more than 5 characters on UNIX-based systems or 61 characters on Windows-based systems. 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. On UNIX-based systems only, you can also choose either the Grid Control or Database Control option if you select Enterprise Manager database management. 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 are on a Windows system or if you select the Database Control option on a UNIX server, then you can set up e-mail notification and enable daily backups. For e-mail notifications, you provide the outgoing mail server and e-mail address. For daily backups, you enter the backup time and operating system credentials for the user that takes backup.

---

**Note:** The option to configure central management for Enterprise Manager is not available during RAC installation on Windows. Also not supported on Windows is the use of standalone Enterprise Manager Configuration Assistant or Database Configuration Assistant to configure central management for RAC. If you want central management for the installed RAC database, then you will have to discover the RAC database target manually from Grid Control after the installation.

---

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 area 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) by checking the select box. If you select a disk with a `Member` status, you must also check the `Force` column to confirm that you want the DBCA to remove the disk from its current ASM 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, note the following

- 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.
- For Windows-based platforms, if the default symbolic links exist and you have not set the `DBCA_RAW_CONFIG` environment variable, then the DBCA replaces the default datafiles with these symbolic link names and displays them in the Database Storage page. If the symbolic links do not exist, then the DBCA displays the default file system datafile file names on the Database Storage page. In this case, replace the default datafile file names with the symbolic link names.

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 Custom 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 Oracle services if you are on a Windows-based platform
- 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:
  - On UNIX-based platforms, execute the dbca command from the \$ORACLE\_HOME/bin directory
  - On Windows-based platforms, choose Start > Programs > Oracle - [HOME\_NAME] > Configuration and Migration Tools > Database Configuration Assistant

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 dialog to confirm the database and instances that the DBCA is going to delete.

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 the selected database's Oracle services for Windows-based platforms
- 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*, *Oracle Database Platform Guide for Windows*, 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:

- [Back Up the Voting Disk after Installation](#)
- [Download and Install Patches](#)
- [Configure Oracle Products](#)
- [Oracle Real Application Clusters 10g Installed on an OCFS](#)
- [Running Oracle9i RAC with Oracle RAC 10g \(Linux Systems Only\)](#)
- [Register COM Applications \(Windows Systems Only\)](#)
- [Ensure Valid Path Name Exists \(Windows Systems Only\)](#)

## Back Up the Voting Disk after Installation

After your Oracle Database 10g with RAC installation is complete and after you are sure that your system is functioning properly, make a backup of the contents of the voting disk. If your installation is on a UNIX system, then use the `dd . exe` utility. If your installation is on a Windows system, then use the command

```
copy <drive>:\cdata\<clustername>\votedsk c:\votedsk.bak
```

if the voting disk is on an Oracle Cluster File System (OCFS) or the following command if is not on an OCFS:

```
ocopy \\.\votedsk c:\votedsk.bak
```

Also make a backup of the voting disk contents after you complete any node additions or node deletions and after running any deinstallation procedures.

## Download and Install Patches

Refer to the Oracle*MetaLink* Web site for required patches for your installation. To download required patches:

1. Use a Web browser to view the Oracle*MetaLink* Web site:

<http://metalink.oracle.com>

2. Log in to Oracle*MetaLink*.

---

---

**Note:** If you are not an Oracle*MetaLink* registered user, then click Register for MetaLink and register.

---

---

3. On the main Oracle*MetaLink* 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 field, 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 Oracle*MetaLink*. On UNIX systems, the unzip utility is located in the \$ORACLE\_HOME/bin directory and on Windows systems, the 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.



## Oracle Real Application Clusters 10g Installed on an OCFS

If you installed RAC on an OCFS, then perform the following steps on each node of the cluster after installation:

- On a UNIX-based system:
  1. Stop the Oracle instance.
  2. Move the `$ORACLE_HOME/dbs/hc_*.dat` files to a directory on a local file system.
  3. Create symbolic links from the `$ORACLE_HOME/dbs` directory to the `hc_*.dat` files on the local file system.
  4. Restart the Oracle instance.
- On a Windows-based system:
  1. Stop the Oracle instance.
  2. Make the files node-specific by executing the command:
 

```
CRS_HOME\cfs\OcfsUtil.exe /c NodeSpecificFile /o create /m o: /p
oracle\product\10.1.0\database\hc_<sid>.dat
```
  3. Restart the Oracle instance.

## Running Oracle9i RAC with Oracle RAC 10g (Linux Systems Only)

If you are running Oracle9i RAC for Linux on the same cluster nodes as Oracle Real Application Clusters 10g, complete the following steps:

---



---

**Note:** These steps are required only if you installed Oracle Real Application Clusters 10g on the same cluster nodes as Oracle9i RAC. If you upgraded from Oracle9i RAC to Oracle Real Application Clusters 10g, do not complete these steps.

---



---

1. Create the following directory:
 

```
$ mkdir -p /etc/ORCLcluster/oracm/lib
```
2. Change directory to this directory:
 

```
$ cd /etc/ORCLcluster/oracm/lib
$
```
3. Copy the `/oracle9i_home/lib/libcmdll.so` file to the current directory:
 

```
$ cp /oracle9i_home/lib/libcmdll.so .
```
4. On any cluster node, enter commands similar to the following to restart the node applications on all cluster nodes:
 

```
$ORACLE_HOME/bin/svrctl stop nodeapps -n nodename
$ORACLE_HOME/bin/svrctl start nodeapps -n nodename
```

In this example, `$ORACLE_HOME` is the Oracle Real Application Clusters 10g Oracle home and `nodename` is the name of the node. Repeat the commands for each node in the cluster.

## Register COM Applications (Windows Systems Only)

On Windows clusters, all COM components must be manually registered on remote cluster nodes. COM components include:

- Oracle Administration Assistant for Windows
- Oracle Counters for Windows Performance Monitor
- Oracle Data Provider for .NET
- Oracle Objects for OLE
- Oracle Provider for OLE DB

To register COM components on remote cluster nodes (the nodes where you did not run the Oracle Universal Installer), execute the following command on each of them:

```
ORACLE_BASE\ORACLE_HOME\bin\selecthome.bat
```

## Ensure Valid Path Name Exists (Windows Systems Only)

Ensure that the path name for your new Oracle home is defined across the cluster by completing the following procedure on each of your cluster nodes:

1. Navigate to **Start > Settings > Control Panel > System > Advanced > Environment Variables**
2. In the System variables dialog, select the **Path** variable and ensure that the value for the **Path** variable contains *Oracle home*\BIN, where Oracle home is your new Oracle home. If the variable does not contain this value, then click **Edit** and add this value to the start of the path variable definition in the Edit System Variable dialog. Click **OK** in the Environment Variables page, then click **OK** in the System Properties page, and then close the Control Panel.
3. Click **OK** in the Environment Variables page, then click **OK** in the System Properties page, and then close the Control Panel.

## Recommended Post-Installation Tasks

This section explains the tasks that Oracle recommends you perform after completing an installation. There is one task recommended on all systems and additional tasks that apply only to UNIX-based or Windows-based systems (but not both), as follows:

- [Verifying Enterprise Manager Operations](#)
- [Recommended Post-Installation Tasks for UNIX](#)
- [Recommended Post-Installation Task for Windows](#)

## Verifying Enterprise Manager Operations

On all systems, you should 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`. 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
```

If you are working on a UNIX-based system, then you should continue with the following section, "[Recommended Post-Installation Tasks for UNIX](#)". If you are working on a Windows-based system, then you should continue with the section, "[Recommended Post-Installation Task for Windows](#)" on page 13-5. When you have completed these platform-specific tasks, you should proceed with the initial configuration tasks described in [Part IV](#).

## Recommended Post-Installation Tasks for UNIX

Oracle recommends that you complete the following tasks after installing Oracle RAC on a UNIX-based system:

- [Back Up the root.sh Script](#)
- [Set Up Users Accounts](#)

### 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*.

When you complete these procedures, you are ready to perform the initial configuration tasks described in [Part IV](#).

## Recommended Post-Installation Task for Windows

Oracle recommends that you complete the following tasks after installing Oracle RAC on a Windows-based system:

- [Configure iSQL\\*Plus for Remote Nodes](#)
- [Using Oracle9i Language and Definition Files with Oracle Database 10g Release 1 \(10.1\)](#)

### Configure iSQL\*Plus for Remote Nodes

If you wish to run *iSQL\*Plus* on a remote node, you need to add a set of directories on each node where you installed a RAC Oracle home. You will not be able to connect to the *iSQL\*Plus* URL prior to creating these directories. To complete this task, add the following directories to the `isqlplus` directory which you can find at *Oracle base\Oracle home\oc4j\j2ee\isqlplus*:

- `connectors`
- `log`
- `persistence`

- `tldcache`

After you create these directories, you can start *iSQL\*Plus* from the Windows Control Panel or by entering `isqlplusctl start` at the command prompt.

### **Using Oracle9i Language and Definition Files with Oracle Database 10g Release 1 (10.1)**

You can use Oracle9i database language and territory definition files with Oracle Database 10g release 1 (10.1).

To enable this functionality:

1. Run the `cr9idata.pl` script, by default located in `ORACLE_BASE\ORACLE_HOME\nls\data\old`.  
  
Alternatively, before you install Oracle Database, you can run the Oracle Universal Installer setup command with the `b_cr9idata` variable set to true, as follows:  

```
setup.exe oracle.rsfnlsrtl_rsf:b_cr9idata=true
```
2. Set the `ORA_NLS10` environment variable to point to the directory where you installed the new language and territory definition files, which by default are in `ORACLE_BASE\ORACLE_HOME\nls\data`.
3. Restart the Oracle database.

# 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 14, "Configuring the Server Parameter File in Real Application Clusters Environments"](#)
- [Chapter 15, "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 platform-dependent.

The default location of the server parameter file on UNIX-based systems is:

```
$ORACLE_HOME/dbs/spfile$ORACLE_SID.ora
```

The default location of the server parameter file on Windows-based platforms is:

```
%ORACLE_HOME%\database\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, for UNIX-based systems Oracle recommends that you use a `PFILE` in this directory:

```
$ORACLE_HOME/dbs/init$ORACLE_SID.ora
```

For Windows-based platforms Oracle recommends that you use a `PFILE` in this directory:

```
%ORACLE_HOME%\database\init%ORACLE_SID%.ora
```

This path is valid for each instance and it refers to a single, shared initialization parameter file. If you use raw storage, then the file must contain the following entry on a UNIX-based system:

```
SPFILE= '/dev/vx/rdisk/oracle_dg/dbspfile'
```

and the following entry on a Windows-based platform:

```
SPFILE= '\\.\dbname_SPFILE'
```

However, if you use a cluster file system, then use one of the following alternate file locations:

- On UNIX-based systems:  

```
SPFILE= '$ORACLE_HOME/dbs/spfile.ora'
```

- On Windows-based systems:  

```
SPFILE=%ORACLE_HOME%\database\spfile.ora
```

If you use ASM, then the `SPFILE` value will be:

```
SPFILE= '+disk_group_name/dbunique_name/spfiledbname.ora'
```

where `dbunique_name` is the unique database name and `dbname` is the database name.



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` on UNIX-based systems and `%ORACLE_HOME%\database\init%ORACLE_SID%.ora` on Windows-based platforms. These are the default `PFILE` names.

---

---

## Parameter File Search Order in Real Application Clusters

Oracle searches for your parameter file in a particular order depending on your platform. On UNIX-based platforms, Oracle examines directories in the following order:

1. `$ORACLE_HOME/dbs/spfilesid.ora`
2. `$ORACLE_HOME/dbs/spfile.ora`
3. `$ORACLE_HOME/dbs/initsid.ora`

The search order on Windows-based platforms is:

1. `%ORACLE_HOME%\database\spfilesid.ora`
2. `%ORACLE_HOME%\database\spfile.ora`
3. `%ORACLE_HOME%\database\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:

---

---

**Note:** The following procedures show examples for UNIX-based systems only.

---

---

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='?/dbs/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/dbs/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](#)
- [UNIX 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 in UNIX-based systems, or a shared logical partition in Windows environments. The Oracle Universal Installer (OUI) automatically initializes the OCR during the CRS installation.

## UNIX 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 15–1](#) shows the tablespace names used by a RAC database and the types of data they contain:

**Table 15–1** 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 <code>DRSYS</code> (contains data for OracleText), <code>CWMLITE</code> (contains the OLAP schemas), <code>XDB</code> (for XML features), <code>ODM</code> (for Oracle Data Mining), <code>TOOLS</code> (contains Enterprise Manager tables), <code>INDEX</code> , <code>EXAMPLE</code> , and <code>OEM-REPO</code> tablespaces. |
| USERS           | Consists of application data. As you create and enter data into tables, Oracle fills this space with your data.                                                                                                                                                                                                                                                                   |

**Table 15–1 (Cont.) Tablespace Names that Real Application Clusters Databases Use**

| Tablespace Name  | Contents                                                                                                                                                                                                                                                                                                                                                          |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 <code>ANALYZE COMPUTE STATISTICS</code> on a very large table, or the constructs <code>GROUP BY</code> , <code>ORDER BY</code> , or <code>DISTINCT</code> . |
| 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 advanced database creation method.

As mentioned, each tablespace has one or more datafiles. The datafile names created by the preconfigured database configuration options vary by operating system and storage type such as ASM, OFS, raw devices, and so on. UNIX-based systems, for example, prompt you to set the file names. Windows-based platforms use the symbolic link names for the datafile and other database files shown in [Table 15–2](#):

**Table 15–2 Windows-Based Platforms Symbolic Link Names and Files**

| Windows-Based Platforms Symbolic Link Names                                                                                   | Tablespace or Other Database Files                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>db_name_system</code>                                                                                                   | SYSTEM                                                                                                                                                                                                                                                                                                                     |
| <code>db_name_SYSAUX</code>                                                                                                   | SYSAUX                                                                                                                                                                                                                                                                                                                     |
| <code>db_name_users</code>                                                                                                    | USERS                                                                                                                                                                                                                                                                                                                      |
| <code>db_name_temp</code>                                                                                                     | TEMP                                                                                                                                                                                                                                                                                                                       |
| <code>db_name_undotbs1</code>                                                                                                 | UNDOTBS1                                                                                                                                                                                                                                                                                                                   |
| <code>db_name_undotbs2</code>                                                                                                 | UNDOTBS2                                                                                                                                                                                                                                                                                                                   |
| <code>db_name_rbs</code>                                                                                                      | RBS (optional)                                                                                                                                                                                                                                                                                                             |
| <code>db_name_example</code>                                                                                                  | EXAMPLE                                                                                                                                                                                                                                                                                                                    |
| <code>db_name_indx</code>                                                                                                     | INDX                                                                                                                                                                                                                                                                                                                       |
| <code>db_name_spfile</code>                                                                                                   | SPFILE                                                                                                                                                                                                                                                                                                                     |
| <code>db_name_control1</code>                                                                                                 | Control File 1                                                                                                                                                                                                                                                                                                             |
| <code>db_name_control2</code>                                                                                                 | Control File 2                                                                                                                                                                                                                                                                                                             |
| <code>db_name_redo_instance_number log_number</code>                                                                          | Redo Log Files                                                                                                                                                                                                                                                                                                             |
| Where <i>instance_number</i> is the number of the instance and <i>log_number</i> is the log number (1 or 2) for the instance. | Each instance must have at least two redo log files. If the database name is <code>db</code> , then the link names for the first instance are:<br><code>db_redo1_1</code><br><code>db_redo1_2</code><br>The link names for the second instance's redo log files are:<br><code>db_redo2_1</code><br><code>db_redo2_2</code> |

You can specify different symbolic names with the Advanced database configuration option.

## 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 advanced 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 14, "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 nondefault 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 advanced 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 Custom Install or Advanced database configuration, 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 15-3](#):

**Table 15–3 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><b>Note:</b> <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> |
| 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>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table 15-3 (Cont.) Connections for Net Service Names**

| Net Service Name Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote listeners      | <p>As discussed in "<a href="#">Configuring Service Registration-Related Parameters in Real Application Clusters</a>" on page 15-4, 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 "<a href="#">Local Listeners</a>" on page 15-5 and "<a href="#">Multiple Listeners</a>" on page 15-6, the <code>LOCAL_LISTENER</code> parameter is set in the <code>init.ora</code> file if a nondefault 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 15–3 (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 15–3 (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 <i>service_name_preconnect</i> 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 = db_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 15–1 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=RO)))
```

**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 Custom Install or Advanced database configuration options 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 binary directory in your CRS home (`CRS home/bin` on UNIX-based systems and `CRS home\BIN` on Windows-based systems) 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](#)

---

---

**Note:** The scripts generated by DBCA are for reference purposes only. Oracle strongly recommends that you use DBCA to create a database.

---

---

## Creating a Database Using Scripts

To generate scripts to create a Real Application Clusters database, create a database using the scripts, 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 scripts 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, which you noted in Step 1, where the DBCA created the scripts, and review the SQL scripts 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. On each cluster node you identified during your DBCA session, execute the script `sid.sh`, on UNIX, or `sid.bat`, on Windows, where *sid* is the sid prefix that you entered on the DBCA Database Name page.
4. Set the initialization parameter, `cluster_database`, to the value TRUE in your SPFILE by issuing an ALTER SYSTEM command, or by uncommenting it in your PFILE for each instance.

5. Configure Net Services to support your new database and instances as described in [Chapter 15, "Understanding the Real Application Clusters Installed Configuration"](#).
6. Set the `local_listener` and `remote_listener` parameters in your SPFILE by issuing an `ALTER SYSTEM` command, or by uncommenting it in your PFILE for each instance.
7. Run SVRCTL to configure and start database and instance applications as described in the *Oracle Real Application Clusters Administrator's Guide*.

---

---

# 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:

- [Raw Devices Required by the DBCA for Non-CFS Environments](#)

## Raw Devices Required by the DBCA for Non-CFS Environments

If you want to use the DBCA to create a database on raw storage, then configure the raw devices described in this section. These devices are in addition to the OCR and voting disk required to install Cluster Ready Services (CRS). 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
- 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 online redo 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 Oracle Database 10g single-instance 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 the same version of Oracle 10g 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 and inventory 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 Enterprise Edition 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 Name [select the database name] > Template Name [enter template name, use database name as the default, and description and template datafile location] > 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 on UNIX and in the `%ORACLE_HOME%\assistants\dbca\templates` directory on Windows-based systems.



## Perform the Pre-Installation Steps

Perform the pre-installation steps as documented in [Part II](#) of this book. For example, on UNIX-based systems, 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 9, "Installing Cluster Ready Services on UNIX"](#) or in [Chapter 10, "Installing Cluster Ready Services on Windows"](#), depending on your operating system.

## 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-2 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 Advanced install type.  
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. On the DBCA Template Selection page, 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.
5. 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 replace default database files with raw devices on the Storage page.

**See Also:** [Chapter 12](#) for more details about the DBCA

6. 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-2. 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 "Extending Clusterware and Oracle Software to New Nodes" 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-2, 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-1 or ["Windows 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 on UNIX or the `OCOPY` command on Windows-based systems. 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 14-3. 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 on UNIX-based systems or `%ORACLE_HOME%\database\initSID.ora` file on Windows-based systems 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$active_instances
```

### Single Instance on a Cluster Running from a RAC-Disabled Oracle Home

On UNIX-based systems, 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-2.
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](#)
- [Windows 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 on UNIX-based systems:

**Table E-1** Directory Structure for A Sample OFA-Compliant UNIX Environment

| Root                       | Second-Level             | Third-Level                  | Fourth-Level                                                                                                    | Fifth-Level                                                                                                      |
|----------------------------|--------------------------|------------------------------|-----------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|
| <code>\$ORACLE_BASE</code> |                          |                              |                                                                                                                 | <code>/u01/app/oracle</code><br>The default <code>ORACLE_BASE</code> directory.                                  |
|                            | <code>ORACLE_HOME</code> |                              |                                                                                                                 | <code>/product/10.1</code><br>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><br><code>/cdump</code><br><code>/hdump</code><br><code>/pfile</code><br><code>/udump</code> | Dump destinations for database server.                                                                           |

**Table E-1 (Cont.) Directory Structure for A Sample OFA-Compliant UNIX Environment**

| Root | Second-Level | Third-Level | Fourth-Level | Fifth-Level                                    |
|------|--------------|-------------|--------------|------------------------------------------------|
|      | CRS Home     |             |              | /crs/10.1<br>The default name of the CRS home. |
|      |              | /bin        |              | Subtree for Oracle binaries.                   |
|      |              | /network    |              | Subtree for Oracle Net.                        |

**See Also:** *Oracle Database 10g Administrator's Reference Release 1 (10.1) for UNIX Systems: AIX-Based Systems, Apple Mac OS X, hp HP-UX, hp Tru64 UNIX, Linux, and Solaris Operating System* for further information about the \$ORACLE\_HOME and /admin directories

## Windows Directory Structures for Real Application Clusters

Table E-2 shows the hierarchical directory tree of a sample OFA-compliant database for RAC on Windows-based systems:

**Table E-2 Directory Structure for A Sample OFA-Compliant Windows Environment**

| Root           | Second-Level   | Third-Level | Fourth-Level | Fifth-Level                                                                                                                                                     |
|----------------|----------------|-------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| x:\oracle_base |                |             |              | c:\oracle<br>The default ORACLE_BASE directory.                                                                                                                 |
|                | \%ORACLE_HOME% |             |              | \ora10.1<br>The name of the Oracle home by default.                                                                                                             |
|                | \%ORA_CRSHOME% |             |              | \crs10.1<br>The name of the CRS home by default.                                                                                                                |
|                |                | \bin        |              | Subtree for Oracle binaries.                                                                                                                                    |
|                |                | \network    |              | Subtree for Oracle Net configuration files, including tnsnames.ora, listener.ora and sqlnet.ora.                                                                |
|                |                | \srvm       |              | \admin subdirectory.                                                                                                                                            |
|                |                |             | \admin       | The RAC script clustdb.sql and initialization parameter files for database creation                                                                             |
|                |                | \database   |              | A directory from previous releases containing initialization files pointing to the new directory location for parameter files: ORACLE_BASE\admin\db_name\pfile. |
|                | \admin         |             |              | Subtree for RAC database administration files                                                                                                                   |
|                |                | \db_name    |              | db_name database administration files for the instance identified by sid.                                                                                       |
|                |                |             | \ad hoc      | Ad hoc SQL scripts.                                                                                                                                             |
|                |                |             | \adump       | Audit files.                                                                                                                                                    |
|                |                |             | \arch        | Archived redo log files.                                                                                                                                        |
|                |                |             | \bdump       | Background process trace files.                                                                                                                                 |
|                |                |             | \cdump       | Core dump files.                                                                                                                                                |
|                |                |             | \create      | Programs used to create the database.                                                                                                                           |



**Table E-2 (Cont.) Directory Structure for A Sample OFA-Compliant Windows**

| <b>Root</b> | <b>Second-Level</b> | <b>Third-Level</b> | <b>Fourth-Level</b> | <b>Fifth-Level</b>             |
|-------------|---------------------|--------------------|---------------------|--------------------------------|
|             |                     |                    | \exp                | Database export files          |
|             |                     |                    | \pfile              | Initialization parameter files |
|             |                     |                    | \udump              | User SQL trace files           |



## Numerics

---

- 32-bit, Windows-based installation, 8-1
- 64-bit
  - checking system architecture, 2-4, 3-3, 5-4, 7-3

## A

---

### Adaptive Server

*See* Sybase Adaptive Server Enterprise

### Additional Real Application Clusters

documentation, 1-1

### Advanced

configuration type, 11-1, 11-2

advanced database type, 11-8

### Advanced Security Option

*See* Oracle Advanced Security

### AdvFS

identifying AdvFS file systems on Tru64 UNIX, 4-31

`aio_task_max_num` subsystem attribute

recommended value on Tru64 UNIX, 4-16

### AIX

activating volume groups, 2-45, 3-41

APAR download location, 2-9

character raw device names, 2-33

checking disk availability for ASM, 2-32, 5-37, 5-40, 5-46, 5-48

checking disk availability for raw devices, 2-41

checking existence of nobody user, 2-13, 4-12, 5-14, 7-14

checking maintenance level, 2-7

checking patch requirements, 2-9

checking software requirements, 2-7

checking version, 2-7

configuring disks for ASM, 2-31

configuring new disks, 2-32, 2-41

configuring raw devices, 2-34

configuring the HAGS socket, 2-47

creating a raw device mapping file, 2-46

creating a volume group, 2-41

creating identical users, 2-14

creating logical volumes, 2-44

creating raw logical volumes, 2-42

creating volume groups, 2-42

identifying disks for ASM, 2-32

identifying disks for LVM, 2-41

identifying LVM disks, 2-32, 2-41

importing raw device disk group, 2-45

initializing disks for LVM, 2-32, 2-41

location of `oratab` file, 2-21, 3-21, 4-19, 5-24, 7-21

patch requirements, 2-7

PTF download location, 2-9

raw device sizes, 2-43

software requirements, 2-6

specifying owner and permissions of raw devices, 2-44, 2-45

verifying that cluster software is running, 2-47

### AIXwindows

Motif library requirement, 2-6

### APAR

checking, 2-9

download location, 2-9

requirements on AIX, 2-7

### architecture

checking system architecture, 2-4, 3-3, 5-4, 7-3

Optimal Flexible Architecture (OFA), 11-3

### archive logs

destinations, conversion to multi-instance, D-2

### ASM

benefits of, 1-7

block and character device names

on HP-UX, 3-29, 3-32

on Solaris, 7-33, 7-36

on Tru64 UNIX, 4-31

block device names on Linux, 5-38, 5-41

changing owner and permissions of disks

on AIX, 2-33

on HP-UX, 3-29, 3-33

on Solaris, 7-33

on Tru64 UNIX, 4-31

changing owner and permissions of raw devices on Linux, 5-43, 5-52

character raw device names on AIX, 2-33

characteristics of failure groups, 1-7, 2-30, 3-26, 4-28, 5-33, 6-26, 7-30, 8-14

checking disk availability on AIX, 2-32, 5-37, 5-40, 5-46, 5-48

checking disk availability on HP-UX, 3-28, 3-31

checking disk availability on Solaris, 7-32, 7-36

checking disk availability on Tru64 UNIX, 4-30

compared to logical volume managers, 1-7

- compared to RAID, 1-7
- configuring disks, 8-16
  - on AIX, 2-31
  - on HP-UX, 3-28
  - on Linux, 5-35
  - on Mac OS X, 6-28
  - on Solaris, 7-32
  - on Tru64 UNIX, 4-30
- disk groups, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
- displaying attached disks
  - on Linux, 5-37, 5-40, 5-46, 5-48
  - on Mac OS X, 6-28
- failure groups, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
  - examples, 2-30, 3-27, 4-29, 5-33, 6-26, 7-30, 8-14
  - identifying, 2-30, 3-27, 4-29, 5-33, 6-26, 7-30, 8-14
- identifying available disks on AIX, 2-32
- identifying available disks on HP-UX, 3-28, 3-31
- identifying available disks on Linux, 5-37, 5-40, 5-46
- identifying available disks on Mac OS X, 6-28
- identifying available disks on Solaris, 7-32, 7-36
- identifying available disks on Tru64 UNIX, 4-30
- identifying disks on AIX, 2-32
- identifying disks on HP-UX, 3-28, 3-31
- identifying disks on Linux, 5-37, 5-40, 5-46
- identifying disks on Mac OS X, 6-28
- identifying disks on Solaris, 7-32, 7-36
- identifying disks on Tru64 UNIX, 4-30
- overview, 1-6
- recommendations for disk groups, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
- redundancy levels, 1-7
- restrictions on use for OCR or voting disk, 2-28, 3-24, 4-26, 5-31, 6-24, 7-28
- space required for preconfigured database, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
- storage option for datafiles, 2-23, 3-23, 4-22, 5-27, 6-23, 7-24, 8-10
- storage option for RAC, 2-24, 3-24, 4-23, 5-27, 6-24, 7-24
- ASM Management Options page
  - Oracle Universal Installer, 11-9
- authentication support
  - pre-installation requirements, 8-21
- authorized problem analysis report
  - See APAR
- automatic undo management, 15-4
- automount
  - enable, 8-9
- Available
  - service configuration policy, 12-2
- Available Product Components page
  - Oracle Universal Installer, 11-7

## B

---

- background process

- crsd, 9-7
- evmd, 9-7
- ocssd, 9-7
- oproc, 9-6
- backups
  - conversion to multi-instance, D-2
- base directory
  - See Oracle base directory
- Bash shell
  - default user startup file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - setting shell limits on Linux, 5-20
- .bash\_profile file, 2-50, 3-45, 4-34, 5-56, 7-46
- Basic
  - TAF failover policy, 12-2
- block device
  - device name on HP-UX, 3-29, 3-32, 3-37
  - device name on Linux, 5-38, 5-41
  - device name on Solaris, 7-33, 7-36, 7-40
  - device name on Tru64 UNIX, 4-31
- Bourne shell
  - default user startup file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - setting shell limits on Linux, 5-20
- bundle
  - checking on HP-UX, 3-7

## C

---

- C compiler
  - requirement on AIX, 2-7
  - requirement on HP-UX, 3-6
  - requirement on Linux, 5-9
  - requirement on Mac OS X, 6-5
  - requirement on Tru64 UNIX, 4-6
  - See also Pro\*C/C++
- C shell
  - default user startup file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - setting shell limits on Linux, 5-20
- certification, hardware and software, 8-5
- cfgmgr command, 2-32, 2-41
- character device
  - device name on HP-UX, 3-29, 3-32, 3-37
  - device name on Solaris, 7-33, 7-36, 7-40
  - device name on Tru64 UNIX, 4-31
- character raw device
  - device name on AIX, 2-33
- chdev command, 2-32, 2-41
- chmod command, 2-22, 2-23, 2-27, 2-33, 3-22, 3-23, 3-29, 3-33, 4-21, 4-22, 4-26, 4-31, 5-25, 5-26, 5-31, 5-43, 5-52, 6-22, 6-23, 7-22, 7-23, 7-28, 7-33
- chown command, 2-22, 2-23, 2-27, 2-33, 3-22, 3-23, 3-29, 3-33, 4-21, 4-22, 4-26, 4-31, 5-25, 5-26, 5-31, 5-43, 5-52, 6-22, 6-23, 7-22, 7-23, 7-28, 7-33
- clit\_stat command, 3-43
- clu\_get\_info command, 4-32
- Cluster Configuration Information page
  - Oracle Universal Installer, 9-3
- Cluster Configuration page

- Oracle Universal Installer, 10-4
- cluster database
  - installed configuration, 15-2
  - reasons not to convert to, D-1
- cluster file system
  - choosing for datafile location, 2-24, 3-24, 4-23, 5-27, 7-24
  - choosing for software location, 2-24, 3-24, 4-23, 5-27, 7-24
  - loading OCFS on Linux, 5-54
  - OCFS on Linux, 5-10
  - requirement for RAC, 2-24, 4-23, 5-28, 7-25
  - storage option for datafiles, 2-23, 3-23, 4-22, 5-27, 7-24, 8-10
  - storage option for RAC, 2-24, 3-24, 4-23, 5-27, 6-24, 7-24
- cluster interconnect
  - Hyper Messaging protocol on HP-UX, 3-6
- cluster name, 9-3, 10-4
- cluster nodes
  - activating volume groups
    - on AIX, 2-45, 3-41
  - importing raw device disk group
    - on AIX, 2-45
    - on HP-UX, 3-40
    - on Solaris, 7-42
  - specifying uids and gids, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - starting on HP-UX, 3-43
- cluster software
  - checking on AIX, 2-47
  - checking on HP-UX, 3-42
  - checking on Linux, 5-53
  - checking on Solaris, 7-44
  - checking on Tru64 UNIX, 4-32
  - starting Serviceguard on HP-UX, 3-43
- clusterware diagnostics, A-1
- cmruncl command, 3-43
- cmrunnode command, 3-43
- cmviewcl command, 3-43
- COM components
  - registration on Windows, 13-4
- compiler
  - checking on Mac OS X, 6-5
- components
  - created when using the DBCA, 15-2
  - system requirements, 8-1
  - Windows system requirements, 8-1
- Configuration Results page
  - VIPCA, 11-11
- configuration types
  - Advanced, 11-1, 11-2
  - Data Warehouse, 11-1
  - Do not create a starter database, 11-1
  - General Purpose, 11-1
  - New Database, 11-1
  - Transaction Processing, 11-1
- Configure Automatic Storage Management page
  - Oracle Universal Installer, 11-10
- configuring
  - global database name, 12-4
  - SID Prefix, 11-8
  - configuring disks for ASM, 8-16
  - configuring raw devices, 8-18
  - connection load balancing, 15-4
  - control file
    - installed configuration, 15-4
  - control files, 1-9
    - described, 15-4
    - raw devices, 8-19
    - raw devices for
      - on AIX, 2-35, 2-43
      - on HP-UX, 3-31, 3-39
      - on Linux, 5-45
      - on Solaris, 7-35, 7-41
  - convert to cluster database
    - administrative considerations, D-2
    - from non-cluster machine, D-2
    - from single-instance, D-4
    - post-conversion, D-7
  - 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, 12-9
  - create database
    - using scripts, B-1
  - creating
    - Real Application Clusters database with the Database Configuration Assistant, 11-3, 12-1, 12-3
  - cross-node registration, 15-4
  - CRS
    - install on UNIX, 9-1
    - install on Windows, 10-1
    - install with Oracle Universal Installer on UNIX, 9-2
    - install with Oracle Universal Installer on Windows, 10-3
    - installation setup procedures for UNIX, 9-1
    - installation setup procedures for Windows, 10-1, 10-3
    - OCR restriction, 2-28, 3-24, 4-26, 5-31, 6-24, 7-28
    - raw device for OCR on AIX, 2-35, 2-43
    - raw device for OCR on HP-UX, 3-31, 3-39, 7-35
    - raw device for OCR on Linux, 5-45
    - raw device for OCR on Solaris, 7-41
    - raw device for voting disk on AIX, 2-35, 2-44, 7-42
    - raw device for voting disk on HP-UX, 3-31, 3-39, 7-35
    - raw device for voting disk on Linux, 5-45
    - voting disk restriction, 2-28, 3-24, 4-26, 5-31, 6-24, 7-28
  - CRS home
    - UNIX-based system, 9-2

- Windows-based system, 10-4
- crsd, 9-7
- CSD
  - download location for MQSeries
    - on Tru64 UNIX, 4-8
  - download location for WebSphere MQ
    - on AIX, 2-9
    - on HP-UX, 3-9
    - on Linux, 5-11
    - on Solaris, 7-11
  - requirements
    - on AIX, 2-8
    - on HP-UX, 3-9
    - on Linux, 5-9
    - on Solaris, 7-10
    - on Tru64 UNIX, 4-8
- csh.login file, 5-21
- csh.login.local file, 5-21
- custom database
  - failure groups for ASM, 2-30, 3-27, 4-29, 5-33, 6-26, 7-30, 8-14
  - requirements when using ASM, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
- Custom installation, 11-7
- Custom installation type
  - reasons for choosing, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11

## D

- data loss
  - minimizing with ASM, 1-7, 2-30, 3-26, 4-28, 5-33, 6-26, 7-30, 8-14
- Data Warehouse
  - configuration type, 11-1, 11-2
- data warehouse database type, 11-8
- database
  - ASM requirements, 1-8
  - components, created when using the DBCA, 15-2
  - configurations, types, 11-2
- Database Configuration Assistant, 11-11
  - ASM Disk Groups page, 12-5
  - components created by, 15-2
  - control files, 15-4
  - Create ASM Instance page, 12-4
  - Create Disk Group page, 12-5
  - creating Real Application Clusters database
    - after installation, 12-3
    - during installation, 11-3, 12-1
  - Creation Options page, 12-9
  - Database Content page, 12-7
  - Database Credentials page, 12-4
  - Database File Locations page, 12-6
  - Database Identification page, 12-3
  - Database Services page, 12-7
  - Database Storage page, 12-8
  - Database Templates page, 12-3
  - datafiles, 15-2
  - deleting databases with, 12-9
  - deleting Real Application Clusters
    - databases, 12-9

- initialization parameter files, 15-4
- Initialization Parameters page, 12-7
- List of Cluster Databases page, 12-9
- Management Options page, 12-4
- Node Selection page, 12-3
- Operations page, 12-3, 12-9
- raw storage requirements, C-1
- Recovery Configuration page, 12-6
- redo log files, 15-4
- rollback segments, 15-4
- Storage Options page, 12-4
- Summary dialog, 12-9, 12-10
- tablespaces, 15-2
  - using, 12-1
- Welcome page, 12-3
- database configuration type
  - selecting, 11-1
- database configuration types, 11-2
- database creation using scripts, B-1
- database directory, E-2
- database domain, 11-8
- database name, 11-8
- databases
  - ASM requirements, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
- datafiles, 1-9
  - and the DBCA, 15-2
  - creating separate directories for, 2-26, 4-25, 5-30, 7-27, 8-12
  - described, 15-2
  - managing with ASM, 1-6
  - minimum disk space for, 2-25, 4-24, 5-29, 7-26, 8-11
  - options for placing on file system, 2-24, 4-23, 5-28, 7-25, 8-10
  - recommendations for file system, 2-25, 4-24, 5-29, 7-25, 8-11
  - setting permissions on datafile directories, 2-27, 4-26, 5-31, 7-28
  - storage options, 2-23, 3-23, 4-22, 5-27, 6-23, 7-24, 8-10
- dba group
  - and ASM disks on AIX, 2-33
  - and ASM disks on HP-UX, 3-29, 3-33
  - and ASM disks on Linux, 5-43, 5-52
  - and ASM disks on Solaris, 7-33
  - and ASM disks on Tru64 UNIX, 4-31
  - and SYSDBA privilege, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
  - creating, 2-11, 3-12, 4-10, 5-13, 6-8, 7-13
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
  - granting privileges on HP-UX, 3-16
  - raw device group
    - on AIX, 2-44, 2-45
    - on HP-UX, 3-40
    - on Linux, 5-51
    - on Solaris, 7-42
- DBCA, 11-11

- delete database, 12-9
- DBCA\_RAW\_CONFIG environment variable, 2-52, 3-46, 5-58, 6-37, 7-48
  - on AIX, 2-39, 2-47, 3-34, 7-39
  - on HP-UX, 3-42
  - on Linux, 5-53
  - on Solaris, 7-44
- DCE
  - patches required on HP-UX, 3-9
  - 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, 15-5
- default file mode creation mask
  - setting, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- de-install CRS software, 11-17
  - UNIX-based systems, 11-17
  - Windows-based systems, 11-18
- de-install Oracle Database software, 11-14
- delete database with DBCA, 12-9
- device names
  - for raw devices on Linux, 5-51
  - IDE disks on Linux, 5-37, 5-41, 5-46
    - on AIX, 2-33
    - on HP-UX, 3-29, 3-32, 3-37
    - on Linux, 5-38, 5-41
    - on Solaris, 7-33, 7-36, 7-40
    - on Tru64 UNIX, 4-31
  - RAID on Linux, 5-38, 5-41, 5-46
  - raw devices on Linux, 5-42, 5-43
  - SCSI disks on Linux, 5-37, 5-41, 5-46
- device numbers
  - identifying major numbers on AIX, 2-41
  - identifying minor numbers on HP-UX, 3-38
- diagnostics, A-1
- directory
  - creating separate datafile directories, 2-26, 4-25, 5-30, 7-27, 8-12
  - database file directory, 2-25, 4-24, 5-29, 7-25, 8-11
  - Oracle base directory, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
  - Oracle home directory, 2-20, 3-20, 4-18, 5-23, 6-20, 7-20
  - Oracle Inventory directory, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
  - oraInventory, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
  - permission for datafile directories, 2-27, 4-26, 5-31, 7-28
- directory structure, E-1
  - UNIX, E-1
  - Windows, E-2
- disk
  - changing permissions and owner for ASM
    - on HP-UX, 3-29, 3-33
- disk automount
  - enable, 8-9
- Disk Configuration - Oracle Cluster Registry (OCR)
  - page
  - Oracle Universal Installer, 10-6
- Disk Configuration - Voting Disk page
  - Oracle Universal Installer, 10-6
- disk devices
  - managing with ASM, 1-6
- disk group
  - ASM, 1-6, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
    - recommendations for ASM disk groups, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
- disk space
  - checking, 2-3, 3-3, 4-3, 5-3, 6-3, 8-5
  - requirement for Oracle base directory, 4-20, 5-24, 5-25, 7-21, 7-22
  - requirements for preconfigured database in
    - ASM, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
- disk space requirements
  - Windows, 8-3
- disklabel command, 4-31
- diskpart
  - automount enable command, 8-9
- diskpart command line tool, 8-16
- disks
  - changing permissions and owner for ASM
    - on AIX, 2-33
    - on HP-UX, 3-29, 3-33
    - on Solaris, 7-33
    - on Tru64 UNIX, 4-31
  - checking availability for ASM on AIX, 2-32, 5-37, 5-40, 5-46, 5-48
  - checking availability for ASM on HP-UX, 3-28, 3-31
  - checking availability for ASM on Solaris, 7-32, 7-36
  - checking availability for ASM on Tru64 UNIX, 4-30
  - checking availability for raw devices on AIX, 2-41
  - checking availability for raw devices on HP-UX, 3-37
  - checking availability for raw devices on Solaris, 7-39
  - configuring for ASM, 8-16
  - configuring for ASM on AIX, 2-31
  - configuring for ASM on HP-UX, 3-28
  - configuring for ASM on Linux, 5-35
  - configuring for ASM on Mac OS X, 6-28
  - configuring for ASM on Solaris, 7-32
  - configuring for ASM on Tru64 UNIX, 4-30
  - configuring new disks, 2-32, 2-41
  - configuring raw devices, 8-18
  - configuring raw devices on AIX, 2-34
  - configuring raw devices on HP-UX, 3-30
  - configuring raw devices on Solaris, 7-34
  - device names on Solaris, 7-40
  - displaying attached disks
    - on Linux, 5-37, 5-40, 5-46, 5-48
    - on Mac OS X, 6-28, 6-31
  - identifying AdvFS file systems on Tru64 UNIX, 4-31
  - identifying file systems on Solaris, 7-33, 7-37, 7-40

- identifying LSM disks on Tru64 UNIX, 4-31
- identifying LVM disks on AIX, 2-32, 2-41
- identifying LVM disks on HP-UX, 3-29, 3-32, 3-37
- identifying LVM disks on Solaris, 7-32, 7-40
- identifying UFS file systems on Tru64 UNIX, 4-31
- identifying unused disks on Tru64 UNIX, 4-31
- raw voting disk on AIX, 2-35, 2-44
- raw voting disk on HP-UX, 3-31, 3-39, 7-35
- raw voting disk on Linux, 5-45
- raw voting disk on Solaris, 7-42
- DISPLAY environment variable, 3-18
  - setting, 2-50, 2-51, 3-44, 3-45, 4-33, 4-34, 5-56, 5-57, 7-46, 7-47
- displaying attached disks
  - on Mac OS X, 6-31
- Distributed Computing Environment
  - See DCE
- Do not create a starter database
  - configuration type, 11-1
- documentation
  - Oracle Real Application Clusters Administrator's Guide, 1-2
  - Oracle Real Application Clusters Deployment and Performance Guide, 1-2
  - Real Application Clusters, 1-1
- dupatch command, 4-8
- Dynamic Server
  - See Informix Dynamic Server

## E

- emulator
  - installing from X emulator, 2-2, 3-2, 4-2, 5-2, 7-2
- enable automount, 8-9
- enable disk automount, 8-9
- End of Installation page
  - Oracle Universal Installer, 9-6
- Enterprise Edition installation, 11-7
- env command, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- environment
  - checking settings, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
  - configuring for oracle user, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- environment variables
  - DBCA\_RAW\_CONFIG, 2-39, 2-47, 3-34, 3-42, 5-53, 7-39, 7-44
  - DISPLAY, 2-50, 2-51, 3-18, 3-44, 3-45, 4-33, 4-34, 5-56, 5-57, 7-46, 7-47
  - ORACLE\_BASE, 2-19, 2-22, 2-23, 2-50, 3-19, 3-22, 3-23, 3-44, 4-18, 4-21, 4-22, 4-33, 5-22, 5-25, 5-26, 5-56, 6-19, 6-22, 6-23, 6-35, 7-19, 7-22, 7-23, 7-46
  - ORACLE\_HOME, 2-49, 2-50, 2-52, 3-44, 3-46, 4-33, 4-35, 5-55, 5-56, 5-58, 6-35, 6-37, 7-45, 7-46, 7-48
  - ORACLE\_SID, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - PATH, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - removing from shell startup file, 2-50, 3-45, 4-34, 5-56, 6-36, 7-47
  - SHELL, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46

- TEMP and TMPDIR, 2-3, 2-51, 3-3, 3-45, 4-3, 4-34, 5-3, 5-57, 6-3, 6-36, 7-3, 7-47, 8-5
- TNS\_ADMIN, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- Environment Variables page
  - Windows, 11-12, 11-13
- equivalence
  - setting up UNIX user equivalence, 2-15, 3-15, 4-14, 7-16, D-3
- errata
  - Linux kernel errata, 5-10
- error messages
  - Real Application Clusters management tools, A-1
- ESQL/C
  - See Informix ESQL/C
  - /etc/csh.login file, 5-21
  - /etc/csh.login.local file, 5-21
  - /etc/hosts.equiv file, 2-15, 3-15, 4-14, 7-16
  - /etc/lvmrc file, 3-36
  - /etc/pam.d/login file, 5-21
  - /etc/privgroup file, 3-16
  - /etc/profile file, 5-21
  - /etc/profile.local file, 5-21
  - /etc/raw file, 5-52
  - /etc/security/limits.so file, 5-20
  - /etc/sysconfig/rawdevices file, 5-42, 5-43, 5-51
  - /etc/sysconfigtab file, 4-16
  - /etc/sysctl.conf file, 5-20, 6-17
  - /etc/system file, 7-18
- evmd, 9-7
- EXAMPLE tablespace
  - raw device for, 8-19
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
- examples
  - ASM failure groups, 2-30, 3-27, 4-29, 5-33, 6-26, 7-30, 8-14
  - Oracle base directories, 2-19, 3-19, 4-17, 5-22, 6-19, 7-19
- exec\_disable\_arg\_limit subsystem attribute
  - recommended value on Tru64 UNIX, 4-15
- external jobs
  - UNIX user required for, 3-10, 4-9, 5-11, 6-6, 7-11
- external redundancy
  - ASM redundancy level, 1-7
- extjob executable
  - UNIX user required for, 3-10, 4-9, 5-11, 6-6, 7-11
- extjob user
  - creating on HP-UX, 3-14

## F

- failover
  - and service registration, 15-4
- failure group
  - ASM, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
  - characteristics of ASM failure group, 1-7, 2-30, 3-26, 4-28, 5-33, 6-26, 7-30, 8-14
  - examples of ASM failure groups, 2-30, 3-27, 4-29,



- 5-33, 6-26, 7-30, 8-14
- fdisk command, 5-37, 5-40, 5-46, 5-48
- features, new, xxv
- fifo\_do\_adaptive subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- file command, 4-30
- file mode creation mask
  - setting, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- file system
  - appropriate for Oracle base directory, 2-22, 3-22, 3-23, 4-21, 5-25, 5-26, 6-21, 6-22, 7-22, 7-23
  - datafile and recovery file placement
    - options, 2-24, 4-23, 5-28, 7-25, 8-10
  - identifying AdvFS file systems on Tru64 UNIX, 4-31
  - identifying formatted devices on Solaris, 7-33, 7-37, 7-40
  - identifying UFS file systems on Tru64 UNIX, 4-31
  - requirements for RAC, 2-24, 4-23, 5-28, 7-25
  - storage option for datafiles, 2-23, 3-23, 4-22, 5-27, 6-23, 7-24, 8-10
  - using for datafiles, 2-25, 4-24, 5-29, 7-25, 8-11
- file systems
  - system requirements, 8-3
  - Windows system requirements, 8-3
- file-max file, 5-19
- file-max parameter
  - recommended value on Linux, 5-19
- files
  - .bash\_profile, 2-50, 3-45, 4-34, 5-56, 7-46
  - control files
    - raw devices, 8-19
    - raw devices for on AIX, 2-35, 2-43
    - raw devices for on HP-UX, 3-31, 3-39
    - raw devices for on Linux, 5-45
    - raw devices for on Solaris, 7-35, 7-41
  - default shell startup file, 5-21
  - editing shell startup file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - /etc/csh.login, 5-21
  - /etc/csh.login.local, 5-21
  - /etc/hosts.equiv, 2-15, 3-15, 4-14, 7-16
  - /etc/lvmrc, 3-36
  - /etc/pam.d/login, 5-21
  - /etc/privgroup, 3-16
  - /etc/profile, 5-21
  - /etc/profile.local, 5-21
  - /etc/raw, 5-52
  - /etc/security/limits.so, 5-20
  - /etc/sysconfig/rawdevices, 5-42, 5-43, 5-51
  - /etc/sysconfigtab, 4-16
  - /etc/sysctl.conf, 5-20, 6-17
  - /etc/system, 7-18
  - .login, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - oraInst.loc, 2-10, 2-20, 3-12, 3-21, 4-10, 4-19, 5-12, 5-23, 6-20, 7-12, 7-20
  - oratab, 2-21, 3-21, 4-19, 5-24, 6-21, 7-21
  - password file
    - raw device for on AIX, 2-35, 2-43
    - raw device for on HP-UX, 3-31, 3-39
    - raw device for on Linux, 5-45
    - raw device for on Solaris, 7-35, 7-41
  - /proc/sys/fs/file-max, 5-19
  - /proc/sys/kernel/sem, 5-19
  - /proc/sys/kernel/shmall, 5-19
  - /proc/sys/kernel/shmmax, 5-19
  - /proc/sys/kernel/shmmni, 5-19
  - /proc/sys/net/ipv4/ip\_local\_port\_range, 5-19
  - .profile, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - profile.local, 5-21
  - raw device mapping file, 2-52, 3-46, 5-58, 6-37, 7-48
    - on AIX, 2-46
    - on HP-UX, 3-41
    - on Linux, 5-52
    - on Solaris, 7-43
    - on Windows, 8-20
  - redo log files
    - raw devices, 8-19
    - raw devices for on AIX, 2-35, 2-43
    - raw devices for on HP-UX, 3-31, 3-39
    - raw devices for on Linux, 5-45
    - raw devices for on Solaris, 7-35, 7-41
  - .rhosts, 2-15, 3-16, 4-14, 7-16
  - server parameter file
    - raw devices, 8-19
    - raw devices for on AIX, 2-35, 2-43
    - raw devices for on HP-UX, 3-31, 3-39
    - raw devices for on Linux, 5-45
    - raw devices for on Solaris, 7-35, 7-41
  - SPFILE
    - raw device for on AIX, 2-35, 2-43
    - raw device for on Solaris, 7-35
    - raw devices, 8-19
    - raw devices for on HP-UX, 3-39
    - raw devices for on Linux, 5-45
    - raw devices for on Solaris, 7-41
  - SPFILE file
    - raw device for on HP-UX, 3-31
- filesets
  - AIX requirements, 2-6
  - checking on AIX, 2-7
  - checking on HP-UX, 3-7
  - for raw devices on Linux, 5-52
  - format command, 7-32, 7-36, 7-39

## G

- Gateway
  - See Oracle Messaging Gateway
- General Purpose
  - configuration type, 11-1
  - description of configuration type, 11-2
- general purpose database type, 11-8
- Generate Database Creation Scripts
  - on Creation Options page, 12-9
- getconf command, 2-4, 3-3, 5-4, 7-3
- getprivgrp command, 3-16
- gh\_chunks subsystem attribute
  - recommended value on Tru64 UNIX, 4-15

- gid
    - identifying existing, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
    - specifying, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
    - specifying on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - global database name, 11-8, 12-4
  - Globally Unique Identifier
    - Windows, 8-4
  - GPT
    - not supported on Windows, 8-4
  - Grid database control, 11-8
  - group IDs
    - identifying existing, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
    - specifying, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
    - specifying on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - groupadd command, 3-12, 4-10, 4-11, 5-13, 7-12, 7-13
  - groups
    - checking for existing oinstall group, 2-10, 3-12, 4-10, 5-12, 6-7, 7-12
    - creating identical groups on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
    - creating the dba group, 2-11, 3-12, 4-10, 5-13, 6-8, 7-13
    - creating the oinstall group, 2-10, 3-11, 4-10, 5-12, 6-7, 7-12
    - creating the oper group, 2-11, 3-12, 4-11, 5-13, 6-9, 7-13
    - specifying when creating UNIX users, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
    - UNIX OSDBA group (dba), 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
    - UNIX OSOPER group (oper), 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
  - GUID Partition Type
    - not supported on Windows, 8-4
- ## H
- 
- HAGS socket
    - configuring on AIX, 2-47
  - hardware certification, 8-5
  - hardware requirements, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
  - high redundancy
    - ASM redundancy level, 1-7
  - home directory
    - See* Oracle home directory
  - hosts.equiv file, 2-15, 3-15, 4-14, 7-16
  - HP 9000 system
    - HP-UX operating system requirement, 3-6
  - HP-UX
    - block and character device names, 3-29, 3-32, 3-37
    - checking disk availability for ASM, 3-28, 3-31
    - checking disk availability for raw devices, 3-37
    - checking HyperFabric software, 3-43
    - checking patch requirements, 3-9
    - checking quality pack requirement, 3-9
    - checking Serviceguard, 3-43
    - checking software requirements, 3-7
    - checking version, 3-7
    - configuring disks for ASM, 3-28
    - configuring kernel parameters, 3-17
    - configuring raw devices, 3-30
    - creating a raw device mapping file, 3-41
    - creating a volume group, 3-37
    - creating a volume group directory on HP-UX, 3-38
    - creating an unprivileged user, 3-14
    - creating logical volumes, 3-39
    - creating raw logical volumes, 3-38
    - creating required X library symbolic links, 3-10
    - creating volume groups, 3-38
    - disabling volume group activation, 3-36
    - granting privileges to OSDBA group, 3-16
    - identifying disks for ASM, 3-28, 3-31
    - identifying disks for LVM, 3-37
    - identifying LVM disks, 3-29, 3-32, 3-37
    - importing raw device disk group, 3-40
    - initializing disks for LVM, 3-37
    - location of oratab file, 2-21, 3-21, 4-19, 5-24, 7-21
    - maxvgs parameter, 3-38
    - patch download location, 3-9
    - patch requirements, 3-8
    - quality pack download location, 3-9
    - quality pack requirements, 3-8
    - raw device sizes, 3-38
    - requirements on HP 9000 systems, 3-6
    - requirements on Itanium systems, 3-6
    - setting kernel parameters, 3-18
    - software requirements, 3-5
    - specifying owner and permissions of raw devices, 3-40
    - starting a cluster node, 3-43
    - System Administration Manager (SAM), 3-18
    - verifying that cluster software is running, 3-42
  - 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, 3-43
    - HP-UX requirement, 3-6
- ## I
- 
- 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, 5-9
    - requirement on Solaris, 7-7
    - requirement on Tru64 UNIX, 4-6
  - id command, 2-13, 2-14, 3-14, 4-12, 4-13, 5-15, 6-11, 6-12, 7-14, 7-15
  - IDE disks
    - device names on Linux, 5-37, 5-41, 5-46
  - Imake

- requirement on HP-UX, 3-6
- importvg command, 2-45
- Informix Dynamic Server
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
- Informix ESQL/C
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
- Ingres II
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
- initialization parameter files, 15-4
  - for instances, 14-1
  - listener parameters, 15-5
- initialization parameters
  - DISPATCHERS, 15-5
  - LOCAL\_LISTENER, 15-5
  - REMOTE\_LISTENER, 15-5, 15-9
- insmod command, 5-54
- installation
  - directory structure, E-1
  - ldap.ora file, 15-7
  - listener.ora file, 15-6
  - overview, 1-5
  - requirements, 8-1
  - requirements, hardware, 1-2
  - requirements, software, 1-2
  - tnsnames.ora file, 15-7
  - verifying raw devices, 12-3
  - Windows requirements, 8-1
- installation log file, 11-6
- installation types
  - and ASM requirements, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
- instance
  - instance identifier (SID), 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - preferences, 12-2
- instance ID, 11-8
- instances
  - initialization parameter files, 14-1
  - SID Prefix, 11-8
- instfix command, 2-9
- ioscan command, 3-28, 3-31, 3-37
- IP Address page
  - VIPCA, 11-11
- ip\_local\_port\_range file, 5-19
- ip\_local\_port\_range parameter
  - recommended value on Linux, 5-19
- ipc
  - kernel subsystem on Tru64 UNIX, 4-15
- ipconfig list
  - Windows setting, 8-8
- isainfo command, 2-4, 3-3, 5-4, 7-3
- Itanium
  - HP-UX operating system requirement, 3-6
- Itanium system
  - HP-UX operating system requirements, 3-6

## J

---

- Java
  - font package requirements for Solaris, 7-7
- Java Runtime Environment
  - requirements, 8-3
  - Windows requirements, 8-3
- JDK
  - AIX requirements, 2-6
  - font packages required on Solaris, 7-7
  - HP-UX requirements, 3-5
  - Tru64 UNIX requirements, 4-5

## K

---

- kernel
  - Linux errata, 5-10
- kernel parameters
  - checking on Linux, 5-19
  - checking on Solaris, 7-18
  - checking on Tru64 UNIX, 4-16
  - configuring on HP-UX, 3-17
  - configuring on Linux, 5-18
  - configuring on Mac OS X, 6-16
  - configuring on Solaris, 7-17
  - configuring on Tru64 UNIX, 4-15
  - making changes persist on Linux, 5-19
  - making changes persist on Solaris, 7-18
  - making changes persist on Tru64 UNIX, 4-16
  - maxvgs on HP-UX, 3-38
  - setting on HP-UX, 3-18
  - setting on Linux, 5-19
- kernel subsystem attributes
  - See* kernel parameters
- Korn shell
  - default user startup file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - setting shell limits on Linux, 5-20
- ksh
  - See* Korn shell
- ksi\_alloc\_max parameter
  - recommended value on HP-UX, 3-17

## L

---

- Language Selection page
  - Oracle Universal Installer, 9-3, 10-4
- ldap.ora file, 15-7
  - creating, 15-7
  - default configuration, 15-7
- library
  - Motif library, 2-6
- Lightweight Directory Access Protocol (LDAP), 15-7
- limit command, 5-21
- limits.so file, 5-20
- Linux
  - binding raw devices to partitions, 5-51
  - block device names, 5-38, 5-41
  - checking distribution, 5-10
  - checking existence of nobody user, 2-13, 4-12, 5-14, 7-14

- checking kernel parameters, 5-19
- checking raw device bindings, 5-42, 5-43
- checking software requirements, 5-10
- checking version, 5-10
- cluster file system, 5-10
- configuring disks for ASM, 5-35
- configuring kernel parameters, 5-18
- creating a raw device mapping file, 5-52
- creating partitions, 5-46
- creating raw partitions, 5-44
- displaying attached disks, 5-37, 5-40, 5-46, 5-48
- IDE disk device names, 5-37, 5-41, 5-46
- identifying disks for ASM, 5-37, 5-40, 5-46
- kernel errata, 5-10
- loading hangcheck-timer module, 5-54
- loading OCFS, 5-54
- location of oratab file, 2-21, 3-21, 4-19, 5-24, 7-21
- making kernel parameter changes persist, 5-19
- RAID device names, 5-38, 5-41, 5-46
- raw device sizes, 5-44
- raw devices, 5-42, 5-43
- SCSI disk device names, 5-37, 5-41, 5-46
- setting kernel parameters, 5-19
- setting shell limits, 5-20
- software requirements, 5-6
- specifying owner and permissions on raw devices, 5-51
- verifying that cluster software is running, 5-53
- List of Cluster Databases page, 12-9
- listener
  - identifying Oracle home for, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45
  - local, 15-4
  - lsnrctl command, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - registration, 15-6
  - stopping, 2-49, 2-50, 3-43, 3-44, 4-32, 4-33, 5-55, 5-56, 6-35, 7-45, 7-46
  - stopping existing listener process, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45, 8-21
- listener.ora file, 15-6
  - configuring, 15-5
  - default configuration, 15-6
- listeners
  - listener.ora file, 15-6
- load balancing
  - and service registration, 15-4
- load\_ocfs command, 5-54
- local database control, 11-8
- local device
  - using for datafiles, 2-25, 4-24, 5-29, 7-26, 8-11
- local listener, 15-4
- log file
  - installation, 11-6
- logical volume manager
  - See LVM
- logical volumes
  - creating on AIX, 2-44
  - creating on HP-UX, 3-39
  - creating on Solaris, 7-42
- .login file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
- login file, 5-21
- lsdev command, 2-32, 2-41, 5-37, 5-40, 5-46, 5-48
- lslpp command, 2-7, 2-9
- LSM
  - identifying volume group devices on Tru64 UNIX, 4-31
- lsmod command, 5-54
- lsnrctl command, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- lspv command, 2-32, 2-41, 2-45
- lvcreate command, 3-39
- LVM
  - checking disk availability on HP-UX, 3-37
  - compared to ASM, 1-7
  - configuring raw devices on AIX, 2-34
  - configuring raw devices on HP-UX, 3-30
  - configuring raw devices on Solaris, 7-34
  - creating a volume group directory on HP-UX, 3-38
  - creating a volume group on AIX, 2-41
  - creating a volume group on HP-UX, 3-37
  - creating a volume group on Solaris, 7-39
  - creating raw logical volumes on AIX, 2-42
  - creating raw logical volumes on HP-UX, 3-38
  - creating raw logical volumes on Solaris, 7-40
  - creating volume groups on AIX, 2-42
  - creating volume groups on HP-UX, 3-38
  - creating volume groups on Solaris, 7-40
  - disabling volume group activation on HP-UX, 3-36
  - identifying available disks on AIX, 2-41
  - identifying available disks on HP-UX, 3-37
  - identifying available disks on Solaris, 7-39
  - identifying LVM disks on HP-UX, 3-37
  - identifying major device numbers on AIX, 2-41
  - identifying minor device numbers on HP-UX, 3-38
  - identifying volume group devices on AIX, 2-32, 2-41
  - identifying volume group devices on HP-UX, 3-29, 3-32
  - identifying volume group devices on Solaris, 7-32, 7-40
  - initializing disks on AIX, 2-32, 2-41
  - initializing disks on HP-UX, 3-37
  - initializing disks on Solaris, 7-40
  - maximum number of volume groups on HP-UX, 3-38
  - recommendations for ASM, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
  - Veritas Volume Manager on Solaris, 7-32, 7-40
- lvmmc file, 3-36

## M

### Mac OS X

- checking operating system version, 6-5
- checking software requirements, 6-5
- configuring disks for ASM, 6-28
- configuring kernel parameters, 6-16
- displaying attached disks, 6-28, 6-31

- identifying disks for ASM, 6-28
- operating system requirements, 6-5
- software requirements, 6-5
- maintenance level
  - checking on AIX, 2-7
- major device numbers
  - identifying on AIX, 2-41
- mapping file
  - for raw devices, 2-52, 3-46, 5-58, 6-37, 7-48
    - on AIX, 2-46
    - on HP-UX, 3-41
    - on Linux, 5-52
    - on Solaris, 7-43
    - on Windows, 8-20
- mask
  - setting default file mode creation mask, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- Master Boot Record partition
  - Windows requirement, 8-4
- max\_async\_req subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- max\_objs subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- 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, 4-16
- max\_thread\_proc parameter
  - recommended value on HP-UX, 3-17
- maxdsiz parameter
  - recommended value on HP-UX, 3-17
- maxdsiz\_64bit parameter
  - recommended value on HP-UX, 3-17
- maxssiz parameter
  - recommended value on HP-UX, 3-17
- maxssiz\_64bit parameter
  - recommended value on HP-UX, 3-17
- maxswapchunks parameter
  - recommended value on HP-UX, 3-17
- maxuprc parameter
  - recommended value on HP-UX, 3-17
- maxvgs parameter, 3-38
- MBR partition
  - Windows requirement, 8-4
- MC/ServiceGuard
  - See* Serviceguard
- memory requirements, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
- Messaging Gateway
  - See* Oracle Messaging Gateway
- MetaLink
  - Oracle patches, 13-2
- migration
  - from single-instance *See* conversion
- minor device numbers
  - identifying on HP-UX, 3-38
- mkdir command, 2-22, 2-23, 2-27, 3-22, 3-23, 4-21,

- 4-22, 4-26, 5-25, 5-26, 5-31, 6-22, 6-23, 7-22, 7-23, 7-28
- mklnv command, 2-44
- mknod command, 3-38
- mkvlg command, 2-42
- mode
  - setting default file mode creation mask, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- Motif
  - requirement on Tru64 UNIX, 4-6
- mount command, 4-31
- mount point
  - for Oracle base directory, 2-19, 3-19, 4-17, 5-21, 6-19, 7-19
- MQSeries
  - checking requirements on Tru64 UNIX, 4-7
  - CSD download location
    - for Tru64 UNIX, 4-8
  - CSDs required
    - on Tru64 UNIX, 4-8
  - requirement on Tru64 UNIX, 4-6
- msg\_size subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- msgmap parameter
  - recommended value on HP-UX, 3-17
- msgmni parameter
  - recommended value on HP-UX, 3-17
- msgseg parameter
  - recommended value on HP-UX, 3-17
- msgsql parameter
  - recommended value on HP-UX, 3-17

## N

---

- NAS
  - requirement for RAC, 2-24, 4-23, 5-28, 7-25
- NCR Teradata ODBC Driver
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
- ncsize parameter
  - recommended value on HP-UX, 3-17
- Net Configuration Assistant, 11-11
- net service names, 15-7
- NetCA, 11-11
- network attached storage
  - See* NAS
- network configuration files
  - ldap.ora, 15-7
  - listener.ora, 15-6
  - sqlnet.ora, 15-12
  - tnsnames.ora, 15-7
- network directory, E-2
- Network Information Services
  - See* NIS
- New Database
  - configuration type, 11-1
- new features, xxv
- new\_wire\_method subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- nfile parameter

- recommended value on HP-UX, 3-17
- nflocks parameter
  - recommended value on HP-UX, 3-17
- ninode parameter
  - recommended value on HP-UX, 3-17
- NIS
  - alternative to local users and groups, 2-10, 3-11, 4-10, 5-12, 6-7, 7-12
- nkthread parameter
  - recommended value on HP-UX, 3-17
- nobody user
  - checking existence of, 2-13, 4-12, 5-14, 6-11, 7-14
  - description, 4-9, 5-11, 6-6, 7-11
- node
  - starting a cluster node on HP-UX, 3-43
- node applic, 11-11
- node applications, 11-11, 12-1
- nofile
  - shell limit on Linux, 5-20
- None
  - TAF failover policy, 12-2
- normal redundancy
  - ASM redundancy level, 1-7
- Not used
  - service configuration policy, 12-2
- nproc
  - shell limit on Linux, 5-20
- nproc parameter
  - recommended value on HP-UX, 3-17
- NTFS
  - system requirements, 8-3

## O

---

- OCFS
  - checking on Linux, 5-10
  - download location for Linux, 5-10
  - installation, 10-7
  - Linux cluster file system, 5-10
  - loading on Linux, 5-54
  - restriction for Oracle base directory on
    - Linux, 2-22, 2-23, 3-22, 3-23, 4-21, 5-25, 5-26, 7-22, 7-23
- OCR
  - installed configuration, 15-1
  - raw device for
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39, 7-35
    - on Linux, 5-45
    - on Solaris, 7-41
  - restriction on using ASM for, 2-28, 3-24, 4-26, 5-31, 6-24, 7-28
  - specify location during install
    - Windows, 10-2
- OCR contents, 15-1
- ocr.loc file, 9-4, 11-17
- ocssd, 9-7
- ODBC driver
  - Teradata requirement on HP-UX, 3-7
  - Teradata requirement on Solaris, 7-8
- OFA
  - recommendations for Oracle base directory, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
  - recommended path for Oracle base
    - directory, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
  - recommended path for Oracle home
    - directory, 2-20, 3-20, 4-18, 5-23, 6-20, 7-20
  - recommended path for Oracle Inventory
    - directory, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
- oinstall group
  - checking for existing, 2-10, 3-12, 4-10, 5-12, 6-7, 7-12
  - creating, 2-10, 3-11, 4-10, 5-12, 6-7, 7-12
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-9, 3-11, 4-9, 5-11, 6-6, 7-11
- olsnodes command, 9-6, A-1
- ONS, 11-11
- Open Client
  - See Sybase Open Client
- Open Database Connectivity
  - See ODBC
- oper group
  - and SYSOPER privilege, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
  - creating, 2-11, 3-12, 4-11, 5-13, 6-9, 7-13
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
- operating system
  - checking distribution and version of Linux, 5-10
  - checking version of AIX, 2-7
  - checking version of HP-UX, 3-7
  - checking version of Solaris, 7-8
  - checking version of Tru64 UNIX, 4-7
- operating system requirements
  - AIX, 2-6
  - HP-UX, 3-5
  - Linux, 5-6
  - Mac OS X, 6-5
  - Solaris, 7-6, 7-7
  - Tru64 UNIX, 4-5
- operating systems, supported, 8-1
- oprocd, 9-6
- Optimal Flexible Architecture
  - See OFA
- Optimal Flexible Architecture (OFA), 11-3
- Oracle Advanced Security
  - patches required on HP-UX, 3-9
  - pre-installation requirements, 8-21
  - requirements on HP-UX, 3-7
  - requirements on Tru64 UNIX, 4-6
- Oracle base directory
  - and ORACLE\_BASE environment variable, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
  - creating, 2-22, 2-23, 3-22, 3-23, 4-21, 4-22, 5-25, 5-26, 6-22, 6-23, 7-22, 7-23
  - creating new, 2-22, 3-22, 4-20, 5-25, 6-21, 7-22
  - description, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
  - determining disk space on, 2-21, 3-22, 4-20, 5-24,

- 6-21, 7-21
- disk space requirements, 4-20, 5-24, 5-25, 7-21, 7-22
- equivalent directory on Windows, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
- examples, 2-19, 3-19, 4-17, 5-22, 6-19, 7-19
- identifying appropriate file system, 2-22, 3-22, 3-23, 4-21, 5-25, 5-26, 6-21, 6-22, 7-22, 7-23
- identifying existing, 2-20, 3-21, 4-19, 5-23, 6-20, 7-20
- mount point for, 2-19, 3-19, 4-17, 5-21, 6-19, 7-19
- OCFS restriction on Linux, 2-22, 2-23, 3-22, 3-23, 4-21, 5-25, 5-26, 7-22, 7-23
- recommended path, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
- relationship with Oracle Software Owner user, 2-19, 3-19, 4-17, 5-22, 6-19, 7-19
- requirement for, 2-18, 3-19, 4-17, 5-21, 6-18, 7-18
- requirement for RAC installations, 2-22, 2-23, 3-22, 3-23, 4-21, 5-25, 5-26, 6-22, 7-22, 7-23
- requirements for existing directory, 2-21, 3-21, 4-20, 5-24, 6-21, 7-21
- Oracle Cluster File System
  - pre-installation steps on Windows, 8-4
  - See* OCFS
- Oracle Cluster Ready Services
  - See* CRS
- Oracle Cluster Registry
  - See* OCR
- Oracle Database
  - creating datafile directories, 2-26, 4-25, 5-30, 7-27, 8-12
  - datafile storage options, 2-23, 3-23, 4-22, 5-27, 6-23, 7-24, 8-10
  - minimum disk space requirements, 2-25, 4-24, 5-29, 7-26, 8-11
  - privileged groups, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
  - requirements with ASM, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
  - setting ORACLE\_SID environment variable, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- Oracle DCE Integration
  - requirements on HP-UX, 3-7
  - requirements on Tru64 UNIX, 4-6
- Oracle Enterprise Manager
  - post-installation configuration, 13-4
  - pre-installation requirements, 8-21
- Oracle Hangcheck Timer
  - loading module on Linux, 5-54
- Oracle home directory
  - description, 2-20, 3-20, 4-18, 5-23, 6-20, 7-20
  - identifying for listener, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45
  - recommended path, 2-20, 3-20, 4-18, 5-23, 6-20, 7-20
  - requirement for, 2-20, 3-20, 4-18, 5-23, 6-20, 7-20
  - requirements, 2-20, 3-20, 4-18, 5-23, 7-20
  - using to identify Oracle base directory, 2-21, 3-21, 4-19, 5-24, 6-21, 7-21
- Oracle home name, 2-20, 3-20, 4-18, 5-23, 6-20, 7-20
- Oracle Inventory
  - description, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
  - pointer file, 2-10, 3-12, 4-10, 5-12, 7-12
- Oracle Inventory directory
  - description, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
  - recommended path, 2-19, 3-19, 4-18, 5-22, 6-19, 7-19
- Oracle Inventory group
  - checking for existing, 2-10, 3-12, 4-10, 5-12, 6-7, 7-12
  - creating, 2-10, 2-11, 3-11, 3-12, 4-10, 5-12, 5-13, 6-7, 7-12
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-9, 3-11, 4-9, 5-11, 6-6, 7-11
- Oracle Messaging Gateway
  - checking requirements on Tru64 UNIX, 4-7
  - CSD requirements
    - on AIX, 2-8
    - on HP-UX, 3-9
    - on Linux, 5-9
    - on Solaris, 7-10
    - on Tru64 UNIX, 4-8
  - requirements on AIX, 2-6
  - requirements on HP-UX, 3-6
  - requirements on Linux, 5-9
  - requirements on Solaris, 7-7
  - requirements on Tru64 UNIX, 4-6
- Oracle Net
  - identifying Oracle home for listener, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45
  - lsnrctl command, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - stopping existing listener, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45, 8-21
  - stopping listener, 2-49, 3-43, 4-32, 5-55, 7-45
  - stopping the listener, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- Oracle Notification Service
  - See* ONS
- Oracle Software Owner user
  - and ASM disks, 2-33, 3-29, 3-33, 4-31, 5-43, 5-52, 7-33
  - configuring environment for, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - creating, 2-12, 3-13, 4-11, 5-13, 6-9, 7-13
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - determining default shell, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - raw device owner
    - on AIX, 2-44, 2-45
    - on HP-UX, 3-40
    - on Linux, 5-51
    - on Solaris, 7-42
  - relationship with Oracle base directory, 2-19, 3-19, 4-17, 5-22, 6-19, 7-19
  - required group membership, 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - setting shell limits for on Linux, 5-20

- setting up user equivalence, 2-15, 3-15, 4-14, 7-16, D-3
- Oracle Spatial
  - Motif requirement on AIX, 2-6
  - Motif requirement on HP-UX, 3-6
  - Motif requirement on Solaris, 7-7
  - Motif requirement on Tru64 UNIX, 4-6
  - requirement for example programs
    - on AIX, 2-6
    - on HP-UX, 3-6
    - on Linux, 5-9
    - on Solaris, 7-7
    - on Tru64 UNIX, 4-6
  - X window requirement on Linux, 5-9
- Oracle Transparent Gateway
  - Informix requirements on HP-UX, 3-7
  - Informix requirements on Solaris, 7-8
  - Ingres requirements on HP-UX, 3-7
  - Ingres requirements on Solaris, 7-8
  - Sybase requirements on AIX, 2-7
  - Sybase requirements on HP-UX, 3-7
  - Sybase requirements on Solaris, 7-8
  - Sybase requirements on Tru64 UNIX, 4-7
  - Teradata requirements on HP-UX, 3-7
  - Teradata requirements on Solaris, 7-8
- Oracle Universal Installer
  - ASM Management Options page, 11-9
  - Available Product Components page, 11-7
  - Cluster Configuration Information page, 9-3
  - Cluster Configuration page, 10-4
  - Configure Automatic Storage Management page, 11-10
  - Disk Configuration - Oracle Cluster Registry (OCR) page, 10-6
  - Disk Configuration - Voting Disk page, 10-6
  - End of Installation page, 9-6
  - Language Selection page, 9-3, 10-4
  - overview of processing, 1-5
  - Privileged Operating System Groups page, 11-10
  - Product-specific Prerequisite Checks page, 11-7
  - Select Data Storage Drive page, 10-5
  - Select Database Configuration page, 11-8
  - Select Database Management Option page, 11-8
  - Select Disk Formatting Options page, 10-4
  - Select Installation Type page, 11-7
  - Select Software Storage Drive page, 10-5
  - Selected Nodes page, 11-6
  - Specify Backup and Recovery Options page, 11-10
  - Specify Database Configuration Options page, 11-8
  - Specify Database File Storage Option page, 11-9
  - Specify Database Schema Passwords page, 11-10
  - Specify File Locations page, 9-3, 10-4, 11-5
  - Specify Hardware Cluster Installation Mode page, 11-6
  - Specify Inventory Directory and Credentials page, 9-3
  - Specify Network Interface Usage page, 9-4, 10-4
  - Summary page, 9-5, 10-6, 11-10
  - Upgrading an Existing Database page, 11-7
  - Voting Disk Information page, 9-4
  - Welcome page, 9-3, 10-4, 11-5, 11-17
- oracle user
  - and ASM disks, 2-33, 3-29, 3-33, 4-31, 5-43, 5-52, 7-33
  - configuring environment for, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - creating, 2-12, 3-13, 4-11, 4-12, 5-13, 5-14, 6-9, 7-13, 7-14
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - determining default shell, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
  - raw device owner
    - on AIX, 2-44, 2-45
    - on HP-UX, 3-40
    - on Linux, 5-51, 7-42
  - relationship with Oracle base directory, 2-19, 3-19, 4-17, 5-22, 6-19, 7-19
  - required group membership, 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - setting shell limits for on Linux, 5-20
  - setting up user equivalence, 2-15, 3-15, 4-14, 7-16, D-3
- ORACLE\_BASE environment variable, 2-19, 2-22, 2-23, 3-19, 3-22, 3-23, 4-18, 4-21, 4-22, 5-22, 5-25, 5-26, 6-19, 6-22, 6-23, 7-19, 7-22, 7-23
- removing from shell startup file, 2-50, 3-45, 4-34, 5-56, 6-36, 7-47
- setting, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- ORACLE\_HOME environment variable
  - removing from shell startup file, 2-50, 3-45, 4-34, 5-56, 6-36, 7-47
  - setting, 2-49, 3-44, 4-33, 5-55, 6-35, 7-45
  - unsetting, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- ORACLE\_SID environment variable
  - removing from shell startup file, 2-50, 3-45, 4-34, 5-56, 6-36, 7-47
  - setting, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- OracleMetaLink, 13-2
- OracleMetaLink site
  - about, 8-5
  - accessing, 8-5
- oraInst.loc file, 2-20, 3-21, 4-19, 5-23, 6-20, 7-20
  - location, 2-10, 3-12, 4-10, 5-12, 7-12
  - location of, 2-10, 3-12, 4-10, 5-12, 7-12
- oraInstRoot.sh script, 9-3
- oraInventory directory
  - See Oracle Inventory directory
- oratab file, 15-2
  - format, 2-21, 3-21, 4-19, 5-24, 6-21, 7-21
  - location of, 2-21, 3-21, 4-19, 5-24, 7-21
- OSDBA group
  - and ASM disks on AIX, 2-33
  - and ASM disks on HP-UX, 3-29, 3-33
  - and ASM disks on Linux, 5-43, 5-52
  - and ASM disks on Solaris, 7-33
  - and ASM disks on Tru64 UNIX, 4-31



- and SYSDBA privilege, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
- creating, 2-11, 3-12, 4-10, 5-13, 6-8, 7-13
- creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
- description, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
- granting privileges on HP-UX, 3-16
- raw device group
  - on AIX, 2-44, 2-45
  - on HP-UX, 3-40
  - on Linux, 5-51
  - on Solaris, 7-42
- oslevel command, 2-7
- OSOPER group
  - and SYSOPER privilege, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
  - creating, 2-11, 3-12, 4-11, 5-13, 6-9, 7-13
  - creating on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - description, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
- OUI
  - See* Oracle Universal Installer

## P

- packages
  - checking on Linux, 5-10
  - checking on Solaris, 7-8
  - requirements for Solaris, 7-7
- parameter file search order, 14-3
- parameters
  - initialization, 14-1
- partition
  - using with ASM, 1-7, 2-28, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
- partitions
  - binding to raw devices on Linux, 5-51
  - creating on Linux, 5-46
  - creating raw partitions on Linux, 5-44
  - required sizes for raw devices on Linux, 5-44
  - slices on Solaris, 7-33, 7-37, 7-40
- partitions, creating with the diskpart command line tool, 8-16
- passwd command, 2-13, 2-14, 3-13, 3-15, 4-12, 4-14, 5-14, 5-16, 6-14, 7-14, 7-16
- password file
  - raw device for
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
- patchadd command, 7-10
- patches
  - checking on AIX, 2-9
  - checking on HP-UX, 3-9
  - checking on Solaris, 7-10
  - download, 13-2
  - download location for AIX, 2-9
  - download location for HP-UX, 3-9
  - download location for Solaris, 7-10
  - download location for Tru64 UNIX, 4-8
  - install, 13-2
  - OracleMetaLink, 13-2
  - requirements on AIX, 2-7
  - requirements on HP-UX, 3-8
  - requirements on Solaris, 7-8
  - requirements on Tru64 UNIX, 4-8
- PATH
  - environment variable
    - set on Windows, 11-12
- PATH environment variable
  - setting, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- PC X server
  - installing from, 2-2, 3-2, 4-2, 5-2, 7-2
- pdisk command, 6-28, 6-31
- 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, 2-27, 4-26, 5-31, 7-28
  - for Oracle base directory, 2-22, 2-23, 3-22, 3-23, 4-21, 4-22, 5-25, 5-26, 6-22, 6-23, 7-22, 7-23
- physical RAM requirements, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
- pkginfo command, 7-8
- post-conversion recommendations, D-7
- post-installation
  - COM components registration, 13-4
  - Oracle Enterprise Manager configuration, 13-4
  - patch download and install, 13-2
  - product configuration, 13-2
  - root.sh back up, 13-5
  - user accounts setup, 13-5
- Precompilers
  - requirements on AIX, 2-7
  - requirements on HP-UX, 3-6
  - requirements on Linux, 5-9
  - requirements on Mac OS X, 6-5
  - requirements on Tru64 UNIX, 4-6
- preconfigured database
  - ASM disk space requirements, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
  - requirements when using ASM, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
- preconfigured database installation types, 11-2
- Pre-connect
  - TAF failover policy, 12-2
- Preferred
  - service configuration policy, 12-2
- pre-installation
  - raw device creation, 15-3, C-1
  - requirements for Oracle Advanced Security, 8-21
  - requirements for Oracle Enterprise Manager, 8-21
- privgroup file, 3-16
- privileged groups
  - for Oracle Database, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
- Privileged Operating System Groups page

- Oracle Universal Installer, 11-10
  - Pro\*C/C++
    - patches required on HP-UX, 3-9
    - patches required on Solaris, 7-10
    - PTFs and APARs required on AIX, 2-8
    - requirements on AIX, 2-7
    - requirements on HP-UX, 3-6
    - requirements on Linux, 5-9
    - requirements on Mac OS X, 6-5
    - requirements on Tru64 UNIX, 4-6
  - proc
    - kernel subsystem on Tru64 UNIX, 4-15
  - process
    - stopping existing, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45
    - stopping existing listener process, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45, 8-21
    - stopping listener process, 2-49, 3-43, 4-32, 5-55, 7-45
  - processor
    - checking system architecture, 2-4, 3-3, 5-4, 7-3
    - /proc/sys/fs/file-max file, 5-19
    - /proc/sys/kernel/sem file, 5-19
    - /proc/sys/kernel/shmall file, 5-19
    - /proc/sys/kernel/shmmni file, 5-19
    - /proc/sys/net/core/rmem\_default file, 5-19
    - /proc/sys/net/core/rmem\_max file, 5-19
    - /proc/sys/net/core/wmem\_default file, 5-19
    - /proc/sys/net/core/wmem\_max file, 5-19
    - /proc/sys/net/ipv4/ip\_local\_port\_range file, 5-19
  - product
    - checking on HP-UX, 3-7
  - Product-specific Prerequisite Checks page
    - Oracle Universal Installer, 11-7
  - .profile file, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - profile file, 5-21
  - profile.local file, 5-21
  - program technical fix
    - See PTF
  - ps command, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45
  - PTF
    - checking, 2-9
    - download location, 2-9
    - requirements on AIX, 2-7
  - Public Network Interfaces page
    - VIPCA, 11-11
  - pvccreate command, 3-37
  - pvdisplay command, 3-29, 3-32, 3-37
- Q**
- 
- quality pack
    - checking requirements on HP-UX, 3-9
    - download location for HP-UX, 3-9
    - requirements on HP-UX, 3-8
- R**
- 
- RAC
    - configuring disks for ASM on HP-UX, 3-28
    - configuring disks for ASM on Linux, 5-37, 5-38, 5-40, 5-41, 5-47, 5-48, 5-49
    - configuring disks for ASM on Solaris, 7-32, 7-36
    - configuring disks for ASM on Tru64 UNIX, 4-30
    - configuring disks for raw devices on HP-UX, 3-37
    - configuring disks for raw devices on Linux, 5-44
    - configuring raw logical volumes on AIX, 2-41
    - configuring raw logical volumes on Solaris, 7-39
    - file system requirements, 2-24, 4-23, 5-28, 7-25
    - patches required on HP-UX, 3-9
    - preferred datafile storage mechanisms, 2-28, 3-24, 4-26, 5-31, 6-24, 7-28
    - requirements for Oracle base directory, 2-22, 2-23, 3-22, 3-23, 4-21, 5-25, 5-26, 6-22, 7-22, 7-23
    - supported storage locations, 2-24, 3-24, 4-23, 5-27, 7-24
    - supported storage options, 2-24, 3-24, 4-23, 5-27, 6-24, 7-24
  - RAC high availability extensions
    - configuring services for, 12-2
    - TAF policies for, 12-2
  - rad\_gh\_regions attribute
    - recommended value on Tru64 UNIX, 4-15
  - RAID
    - compared to ASM, 1-7
    - device names on Linux, 5-38, 5-41, 5-46
    - recommended ASM redundancy level, 1-7, 2-29, 3-25, 4-27, 5-32, 6-25, 7-29, 8-13
    - using for Oracle datafiles, 2-25, 4-24, 5-29, 7-25, 8-11
  - RAM requirements, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
  - raw command, 5-42, 5-43, 5-51
  - raw device
    - for OCR
      - on AIX, 2-35, 2-43
      - on HP-UX, 3-31, 3-39, 7-35
      - on Linux, 5-45
      - on Solaris, 7-41
    - for password file
      - on AIX, 2-35, 2-43
      - on HP-UX, 3-31, 3-39
      - on Linux, 5-45
      - on Solaris, 7-35, 7-41
    - for server parameter file, 8-19
      - on AIX, 2-35, 2-43, 3-31
      - on HP-UX, 3-39
      - on Linux, 5-45
      - on Solaris, 7-35, 7-41
    - for SPFILE, 8-19
      - on AIX, 2-35, 2-43
      - on HP-UX, 3-31, 3-39
      - on Linux, 5-45
      - on Solaris, 7-35, 7-41
    - for voting disk
      - on AIX, 2-35, 2-44
      - on HP-UX, 3-31, 3-39, 7-35
      - on Linux, 5-45
      - on Solaris, 7-42
  - raw devices
    - activating volume group on cluster nodes

- on AIX, 2-45, 3-41
- binding on Linux, 5-42, 5-43
- binding to partitions on Linux, 5-51
- block and character device names on HP-UX, 3-37
- block and character device names on Solaris, 7-40
- changing permissions and owner for ASM on Linux, 5-43, 5-52
- checking availability of disks on AIX, 2-41
- checking disk availability on HP-UX, 3-37
- checking disk availability on Solaris, 7-39
- configuring, 8-18
- configuring on AIX, 2-34
- configuring on HP-UX, 3-30
- configuring on Solaris, 7-34
- creating a raw device mapping file, 6-33
  - on AIX, 2-46
  - on HP-UX, 3-41
  - on Linux, 5-52
  - on Solaris, 7-43
  - on Windows, 8-20
- creating logical volumes on AIX, 2-44
- creating logical volumes on HP-UX, 3-39
- creating logical volumes on Solaris, 7-42
- creating partitions on Linux, 5-46
- creating raw logical volumes on AIX, 2-42
- creating raw logical volumes on HP-UX, 3-38
- creating raw logical volumes on Solaris, 7-40
- creating raw partitions on Linux, 5-44
- Database Configuration Assistant, C-1
- DBCA\_RAW\_CONFIG environment variable, 2-52, 3-46, 5-58, 6-37, 7-48
- device names on Linux, 5-42, 5-43, 5-51
  - device names, 5-52
- for control files, 8-19
  - on AIX, 2-35, 2-43, 3-31, 7-35
  - on HP-UX, 3-39
  - on Linux, 5-45
  - on Solaris, 7-41
- for EXAMPLE tablespace, 8-19
  - on AIX, 2-35, 2-43
  - on HP-UX, 3-31, 3-39
  - on Linux, 5-45
  - on Solaris, 7-35, 7-41
- for redo log files, 8-19
  - on AIX, 2-35, 2-43, 3-31
  - on HP-UX, 3-39
  - on Linux, 5-45
  - on Solaris, 7-35, 7-41
- for SYSAUX tablespace, 8-19
  - on AIX, 2-34, 2-43
  - on HP-UX, 3-30, 3-39
  - on Linux, 5-44
  - on Solaris, 7-35, 7-41
- for SYSTEM tablespace, 8-19
  - on AIX, 2-34, 2-43
  - on HP-UX, 3-30, 3-38
  - on Linux, 5-44
  - on Solaris, 7-35, 7-41
- for TEMP tablespace, 8-19
  - on AIX, 2-35, 2-43
  - on HP-UX, 3-31, 3-39
  - on Linux, 5-45
  - on Solaris, 7-35, 7-41
- for UNDOTBS tablespace, 8-19
  - on AIX, 2-35, 2-43
  - on HP-UX, 3-30, 3-39
  - on Linux, 5-45
  - on Solaris, 7-35, 7-41
- for USER tablespace
  - on AIX, 2-35, 2-43
  - on HP-UX, 3-31, 3-39
  - on Linux, 5-45
  - on Solaris, 7-35, 7-41
- for USERS tablespace, 8-19
- identifying disks on AIX, 2-41
- identifying disks on HP-UX, 3-37
- identifying disks on Solaris, 7-39
- identifying LVM disks on HP-UX, 3-37
- identifying on Windows
  - db\_name\_indx1 symbolic link, 15-3
  - db\_name\_rbs1 symbolic link, 15-3
  - db\_name\_redo\_thread\_number symbolic link, 15-3
  - db\_name\_spfile1 symbolic link, 15-3
  - db\_name\_system1 symbolic link, 15-3
  - db\_name\_temp1 symbolic link, 15-3
  - db\_name\_users1 symbolic link, 15-3
- identifying on Windows, undotbs1 and undotbs2 symbolic links, 15-3
- importing on disk group on cluster nodes
  - on AIX, 2-45
  - on HP-UX, 3-40
  - on Solaris, 7-42
- initializing disks for LVM on AIX, 2-32, 2-41
- initializing disks for LVM on HP-UX, 3-37
- initializing disks for LVM on Solaris, 7-40
- required sizes
  - on AIX, 2-43
  - on HP-UX, 3-38
  - on Linux, 5-44
  - on Solaris, 7-41
- setting up, 15-3, C-1
- specifying location of mapping file, 2-52, 3-46, 5-58, 6-37, 7-48
- specifying owner and permissions on AIX, 2-44, 2-45
- specifying owner and permissions on HP-UX, 3-40
- specifying owner and permissions on Linux, 5-51
- specifying owner and permissions on Solaris, 7-42
- storage option for datafiles, 2-23, 3-23, 4-22, 5-27, 6-23, 7-24, 8-10
- storage option for RAC, 2-24, 3-24, 4-23, 5-27, 6-24, 7-24
- value for DBCA\_RAW\_CONFIG environment variable
  - on AIX, 2-39, 2-47, 3-34, 7-39
  - on HP-UX, 3-42

- on Linux, 5-53
  - on Solaris, 7-44
- verification, 12-3
- raw file, 5-52
- raw logical volumes
  - See raw devices
- raw partitions
  - configuring on Solaris, 7-34
  - create on Windows, 8-4
  - See raw devices
- raw storage
  - Database Configuration Assistant
    - requirements, C-1
- rawdevices file, 5-42, 5-43, 5-51
- RBS tablespace
  - description, 15-3
- rcp command, 2-15, 3-15, 4-14, 7-16
- rdg
  - kernel subsystem on Tru64 UNIX, 4-16
- rdg\_auto\_msg\_wires subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- rdg\_max\_auto\_msg\_wires subsystem attribute
  - recommended value on Tru64 UNIX, 4-16
- Real Application Clusters
  - components, 1-9
  - databases, deleting, 12-9
  - installation requirements, 1-2
  - installed components of, 1-9
  - management tools error messages, A-1
  - overview, 1-1, 12-1, 15-1
  - raw device setup, 15-3, C-1
  - See RAC
- reboot command, 7-18
- recovery files
  - options for placing on file system, 2-24, 4-23, 5-28, 7-25, 8-10
- Red Hat
  - binding partitions to raw devices, 5-51
  - operating system requirements, 5-6, 5-7, 5-9
- Red Hat Package Manager
  - See RPM
- redo log file
  - installed configuration, 15-4
- redo log files, 1-10
  - described, 15-4
  - raw devices for, 8-19
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
- redundancy level
  - and space requirements for preconfigured
    - database, 1-8, 2-29, 3-26, 4-28, 5-33, 6-25, 7-30, 8-14
  - for ASM, 1-7
- redundant array of independant disks
  - See RAID
- redundant array of independent disks
  - See RAID
- registration

- cross-node, 15-4
- remote shell
  - configuring user equivalence for, 2-15, 3-15, 4-14, 7-16
- remsh command, 3-16
- Rendezvous
  - requirement on HP-UX, 3-6
  - requirement on Linux, 5-9
  - requirement on Solaris, 7-7
- required service packs, 8-1
- requirements
  - for installation, 8-1
  - for installation on Windows, 8-1
  - for Java Runtime Environment, 8-3
  - for Java Runtime Environment on Windows, 8-3
  - for Oracle Enterprise Manager, 8-21
  - hardware, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
  - hardware certification, 8-5
  - software certification, 8-5
  - Windows Remote Desktop Connection
    - support, 8-6
  - Windows service packs, 8-2
  - Windows web browser support, 8-5
- .rhosts file, 2-15, 3-16, 4-14, 7-16
- rmem\_default file, 5-19
- rmem\_default parameter
  - recommended value on Linux, 5-19
- rmem\_max file, 5-19
- rmem\_max parameter
  - recommended value on Linux, 5-19
- rollback segments
  - described, 15-4
- root user
  - logging in as, 2-2, 3-1, 4-1, 5-2, 6-1, 7-2
- root.sh, 11-11
  - back up, 13-5
- RPM
  - checking on Linux, 5-10
- rpm command, 5-10
- rsh command, 2-15, 4-14, 7-17
  - configuring user equivalence for, 2-15, 3-15, 4-14, 7-16
- rt
  - kernel subsystem on Tru64 UNIX, 4-16

## S

- SAM
  - starting on HP-UX, 3-18
- sam command, 3-18
- Save as a Database Template
  - option on Creation options page, 12-9
- scp command, 2-15, 3-15, 4-14, 7-16
- scripts to create a Real Application Clusters
  - database, B-1
- SCSI disks
  - device names on Linux, 5-37, 5-41, 5-46
- scstat command, 7-45
- secure shell
  - as alternative to remote shell (rsh), 2-15, 3-15,

- 4-14, 7-16
- Select Data Storage Drive page
  - Oracle Universal Installer, 10-5
- Select Database Configuration page
  - Oracle Universal Installer, 11-8
- Select Database Management Option page
  - Oracle Universal Installer, 11-8
- Select Disk Formatting Options page
  - Oracle Universal Installer, 10-4
- Select Installation Type page
  - Oracle Universal Installer, 11-7
- Select Software Storage Drive page
  - Oracle Universal Installer, 10-5
- Selected Nodes page
  - Oracle Universal Installer, 11-6
- sem file, 5-19
- seminfo\_semgni parameter
  - recommended value on Solaris, 7-17
- seminfo\_semmns parameter
  - recommended value on Solaris, 7-17
- seminfo\_semmnsl parameter
  - recommended value on Solaris, 7-17
- seminfo\_semvmx parameter
  - recommended value on Solaris, 7-17
- semmap parameter
  - recommended value on HP-UX, 3-17
- semgni parameter
  - recommended value on HP-UX, 3-17
  - recommended value on Linux, 5-19
  - recommended value on Solaris, 7-17
- semmns parameter
  - recommended value on HP-UX, 3-17
  - recommended value on Linux, 5-19
  - recommended value on Solaris, 7-17
- semmnu parameter
  - recommended value on HP-UX, 3-17
- semmnsl parameter
  - recommended value on Linux, 5-19
  - recommended value on Solaris, 7-17
- semopm parameter
  - recommended value on Linux, 5-19
- semume parameter
  - recommended value on Mac OS X, 6-16
- semvmx parameter
  - recommended value on HP-UX, 3-17
  - recommended value on Solaris, 7-17
- server parameter file
  - raw device, 8-19
  - raw device for
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
- server parameter files, 1-9, 14-1, 15-4
  - errors, 14-4
- service command, 5-43, 5-51
- Service Management
  - using, 12-2
- service pack requirements
  - Windows, 8-2
- service packs, 8-1
- service registration
  - configuring, 15-4
- Serviceguard
  - checking on HP-UX, 3-43
  - HP-UX requirement, 3-6
  - starting a cluster node, 3-43
  - starting on HP-UX, 3-43
- services, 12-2
  - stopping existing, 8-21
- setld command, 4-7
- setprivgrp command, 3-16
- shared configuration file, 15-1
- shared server, 15-5
- shared storage
  - requirement for RAC, 2-24, 3-24, 4-23, 5-27, 7-24
- shell
  - default shell startup file, 5-21
  - determining default shell for oracle user, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- SHELL environment variable
  - checking value of, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- shell limits
  - setting on Linux, 5-20
- shell startup file
  - editing, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
  - removing environment variables, 2-50, 3-45, 4-34, 5-56, 6-36, 7-47
- 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, 5-19
- shmall parameter
  - recommended value on Linux, 5-19
- shminfo\_shmmax parameter
  - recommended value on Solaris, 7-17
- shminfo\_shmmin parameter
  - recommended value on Solaris, 7-17
- shminfo\_shmgni parameter
  - recommended value on Solaris, 7-17
- shminfo\_shmseg parameter
  - recommended value on Solaris, 7-17
- shmmax file, 5-19
- shmmax parameter
  - recommended value on HP-UX, 3-17
  - recommended value on Linux, 5-19
  - recommended value on Solaris, 7-17
- shmmin parameter
  - recommended value on Solaris, 7-17
- shmmni file, 5-19
- shmmni parameter
  - recommended value on HP-UX, 3-17
  - recommended value on Linux, 5-19
  - recommended value on Solaris, 7-17
- shmseg parameter

- recommended value on HP-UX, 3-17
- recommended value on Solaris, 7-17
- shutdown command, 4-17
- SID, 11-8
  - setting ORACLE\_SID environment variable, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- sid Prefix, 11-8
- sizer command, 4-7
- smit command, 2-11, 2-12, 2-14
- software certification, 8-5
- Software Only
  - configuration type, 11-2
- software requirements
  - AIX, 2-6
  - checking on AIX, 2-7
  - checking on HP-UX, 3-7
  - checking on Linux, 5-10
  - checking on Solaris, 7-8
  - checking on Tru64 UNIX, 4-7
  - HP-UX, 3-5
  - Linux, 5-6
  - Mac OS X, 6-5
  - Solaris, 7-6
  - Tru64 UNIX, 4-5
- Solaris
  - block and character device names, 7-33, 7-36, 7-40
  - checking disk availability for ASM, 7-32, 7-36
  - checking disk availability for raw devices, 7-39
  - checking existence of nobody user, 2-13, 4-12, 5-14, 7-14
  - checking kernel parameters, 7-18
  - checking patch requirements, 7-10
  - checking software requirements, 7-8
  - checking Sun Cluster software, 7-45
  - checking version, 7-8
  - configuring disks for ASM, 7-32
  - configuring kernel parameters, 7-17
  - configuring raw devices, 7-34
  - configuring raw partitions, 7-34
  - creating a raw device mapping file, 7-43
  - creating a volume group, 7-39
  - creating logical volumes, 7-42
  - creating raw logical volumes, 7-40
  - creating volume groups, 7-40
  - font packages for Java, 7-7
  - identifying disks for ASM, 7-32, 7-36
  - identifying disks for LVM, 7-39
  - identifying file systems, 7-33, 7-37, 7-40
  - identifying LVM disks, 7-32, 7-40
  - importing raw device disk group, 7-42
  - initializing disks for LVM, 7-40
  - location of oratab file, 2-21, 3-21, 4-19, 5-24, 7-21
  - making kernel parameter changes persist, 7-18
  - operating system requirements, 7-7
  - package requirements, 7-7
  - patch download location, 7-10
  - patch requirements, 7-8
  - raw device sizes, 7-41
  - software requirements, 7-6
  - specifying owner and permissions on raw devices, 7-42
  - Sun Cluster requirement, 7-7
  - verifying that cluster software is running, 7-44
  - Veritas Volume Manager, 7-32, 7-40
- Specif File Locations page
  - Oracle Universal Installer, 9-3
- Specify Backup and Recovery Options page
  - Oracle Universal Installer, 11-10
- Specify Database Configuration Options page
  - Oracle Universal Installer, 11-8
- Specify Database File Storage Option page
  - Oracle Universal Installer, 11-9
- Specify Database Schema Passwords page
  - Oracle Universal Installer, 11-10
- Specify File Locations page
  - Oracle Universal Installer, 10-4
  - Specify File Locations page, 11-5
- Specify Hardware Cluster Installation Mode page
  - Oracle Universal Installer, 11-6
- Specify Inventory Directory and Credentials page
  - Oracle Universal Installer, 9-3
- Specify Network Interface Usage page
  - Oracle Universal Installer, 9-4, 10-4
- SPFILE
  - default creation, 14-1
  - default location, 14-2
  - managing on raw device, 14-2
  - migrate to, 14-3
  - raw device for
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
  - raw devices, 8-19
  - sqlnet.ora file, 15-12
    - default configuration, 15-12
  - srvm\admin directory, E-2
- SSH
  - as alternative to remote shell (rsh), 2-15, 3-15, 4-14, 7-16
- ssm\_threshold subsystem attribute
  - recommended value on Tru64 UNIX, 4-15
- Standard Edition installation, 11-7
- startup file
  - default shell startup file, 5-21
  - for shell, 2-50, 3-45, 4-34, 5-56, 6-35, 7-46
- stopping
  - existing services, 8-21
- storage device
  - shared storage requirements, 2-24, 3-24, 4-23, 5-27, 7-24
- subsets
  - checking on Tru64 UNIX, 4-7
  - requirements for Tru64 UNIX, 4-6
- Summary dialog, 12-9, 12-10
- Summary page
  - Oracle Universal Installer, 9-5, 10-6, 11-10
  - VIPCA, 11-11
- Sun Cluster
  - checking on Solaris, 7-45

- patches required on Solaris, 7-10
- required UDLM patch, 9-2
- requirement on Solaris, 7-7
- supported operating systems, 8-1
- SuSE
  - binding partitions to raw devices, 5-51
  - operating system requirements, 5-6, 5-7, 5-9
- swap space
  - requirements, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
- swlist command, 3-7, 3-9
- Sybase Adaptive Server Enterprise
  - requirements on AIX, 2-7
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
  - requirements on Tru64 UNIX, 4-7
- Sybase Open Client
  - requirements on AIX, 2-7
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
  - requirements on Tru64 UNIX, 4-7
- symbolic links
  - X library links required on HP-UX, 3-10
- SYSAUX tablespace
  - for raw devices, 8-19
  - raw device for
    - on AIX, 2-34, 2-43
    - on HP-UX, 3-30, 3-39
    - on Linux, 5-44
    - on Solaris, 7-35, 7-41
- sysconfig command, 4-16
- sysconfigdb command, 4-17
- sysconfigtab file, 4-16
- sysctl command, 5-19, 6-17
- sysctl.conf file, 5-20, 6-17
- SYSDBA privilege
  - associated UNIX group, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
- sysdef command, 7-18
- SYSOPER privilege
  - associated UNIX group, 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
- System Administration Manager
  - See* SAM
- system architecture
  - checking, 2-4, 3-3, 5-4, 7-3
- system file, 7-18
- system requirements, 8-1
  - for NTFS file systems, 8-3
  - for Windows file systems, 8-3
  - Windows, 8-1
- SYSTEM tablespace
  - description, 15-2
  - for raw devices, 8-19
  - raw device for
    - on AIX, 2-34, 2-43
    - on HP-UX, 3-30, 3-38
    - on Linux, 5-44
    - on Solaris, 7-35, 7-41
- system variables dialog
  - Windows, 11-12, 11-13

## T

---

- tablespaces
  - and the DBCA, 15-2
  - expanding for large sorts, 15-3
  - RBS, 15-3
  - SYSTEM, 15-2
  - TEMP, 15-3
  - undo tablespaces for automatic undo management, 15-3
  - USERS, 15-2
- tcsh shell
  - setting shell limits on Linux, 5-20
- TEMP environment variable, 2-3, 3-3, 4-3, 5-3, 6-3, 7-3, 8-5
  - setting, 2-51, 3-45, 4-34, 5-57, 6-36, 7-47
- TEMP tablespace
  - described, 15-3
  - for raw devices, 8-19
  - raw device for
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
- temporary directory, 2-3, 3-3, 4-3, 5-3, 6-2, 7-3
- temporary directory on Windows, 8-5
- temporary disk space
  - checking, 2-3, 3-3, 4-3, 5-3, 6-2, 7-3
  - checking on Windows, 8-5
  - freeing, 2-3, 3-3, 4-3, 5-3, 6-2, 7-3
  - freeing on Windows, 8-5
  - requirements, 2-3, 3-2, 4-2, 5-3, 6-2, 7-3
- Teradata ODBC Driver
  - requirements on HP-UX, 3-7
  - requirements on Solaris, 7-8
- TIBCO Rendezvous
  - requirement on HP-UX, 3-6
  - requirement on Linux, 5-9
  - requirement on Solaris, 7-7
- /tmp directory
  - checking space in, 2-3, 3-3, 4-3, 5-3, 6-2, 7-3
  - freeing space in, 2-3, 3-3, 4-3, 5-3, 6-2, 7-3
- tmp directory
  - checking space for Windows installation, 8-5
  - freeing space for Windows installation, 8-5
- TMPDIR environment variable, 2-3, 3-3, 4-3, 5-3, 6-3, 7-3, 8-5
  - setting, 2-51, 3-45, 4-34, 5-57, 6-36, 7-47
- TNS\_ADMIN environment variable
  - unsetting, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- tnsnames.ora file, 15-7
  - default configuration, 15-7
- Transaction Processing
  - configuration type, 11-1
  - description of configuration type, 11-2
  - transaction processing database type, 11-8
- Transparent Application Failover (TAF)
  - policies, 12-2
- Transparent Gateway
  - See* Oracle Transparent Gateway
- Tru64 UNIX

- block and character device names, 4-31
- checking disk availability for ASM, 4-30
- checking existence of nobody user, 2-13, 4-12, 5-14, 7-14
- checking kernel subsystem attributes, 4-16
- checking software requirements, 4-7
- checking version, 4-7
- configuring disks for ASM, 4-30
- configuring kernel parameters, 4-15
- identifying AdvFS file systems, 4-31
- identifying disks for ASM, 4-30
- identifying LSM disks, 4-31
- identifying UFS file systems, 4-31
- identifying unused disks, 4-31
- location of oratab file, 2-21, 3-21, 4-19, 5-24, 7-21
- making subsystem attribute value changes
  - persist, 4-16
- patch download location, 4-8
- patch requirements, 4-8
- software requirements, 4-5
- subset requirements, 4-6
- TruCluster requirement, 4-6
- verifying that cluster software is running, 4-32

TruCluster

- requirement for Tru64 UNIX, 4-6

## U

### UDLM

- required patch for Sun Clusters, 9-2

### UFS

- identifying UFS file systems on Tru64 UNIX, 4-31

### uid

- identifying existing, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
- specifying, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
- specifying on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15

ulimit command, 5-21

umask, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48

umask command, 2-50, 2-52, 3-44, 3-46, 4-33, 4-35, 5-56, 5-58, 6-35, 6-37, 7-46, 7-48

uname command, 3-7, 7-8

undo management, 15-4

### UNDOTBS tablespace

- raw device, 8-19

- raw device for

- on AIX, 2-35, 2-43
- on HP-UX, 3-30, 3-39
- on Linux, 5-45
- on Solaris, 7-35, 7-41

UNIX commands, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48

- cfgmgr, 2-32, 2-41

- chdev, 2-32, 2-41

- chmod, 2-22, 2-23, 2-27, 2-33, 3-22, 3-23, 3-29, 3-33, 4-21, 4-22, 4-26, 4-31, 5-25, 5-26, 5-31, 5-43, 5-52, 6-22, 6-23, 7-22, 7-23, 7-28, 7-33

- chown, 2-22, 2-23, 2-27, 2-33, 3-22, 3-23, 3-29, 3-33, 4-21, 4-22, 4-26, 4-31, 5-25, 5-26, 5-31, 5-43, 5-52, 6-22, 6-23, 7-22, 7-23, 7-28, 7-33

clic\_stat, 3-43

clu\_get\_info, 4-32

cmruncl, 3-43

cmrunnode, 3-43

cmviewcl, 3-43

disklabel, 4-31

dupatch, 4-8

env, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48

fdisk, 5-37, 5-40, 5-46, 5-48

file, 4-30

format, 7-32, 7-36, 7-39

getconf, 2-4, 3-3, 5-4, 7-3

getprivgrp, 3-16

groupadd, 3-12, 4-10, 4-11, 5-13, 7-12, 7-13

id, 2-13, 2-14, 3-14, 4-12, 4-13, 5-15, 6-11, 6-12, 7-14, 7-15

importvg, 2-45

insmod, 5-54

instfix, 2-9

ioscan, 3-28, 3-31, 3-37

isainfo, 2-4, 3-3, 5-4, 7-3

limit, 5-21

load\_ocfs, 5-54

lsdev, 2-32, 2-41, 5-37, 5-40, 5-46, 5-48

lslpp, 2-7, 2-9

lsmod, 5-54

lspv, 2-32, 2-41, 2-45

lvcreate, 3-39

mkdir, 2-22, 2-23, 2-27, 3-22, 3-23, 4-21, 4-22, 4-26, 5-25, 5-26, 5-31, 6-22, 6-23, 7-22, 7-23, 7-28

mkiv, 2-44

mknod, 3-38

mkvg, 2-42

mount, 4-31

oslevel, 2-7

passwd, 2-13, 2-14, 3-13, 3-15, 4-12, 4-14, 5-14, 5-16, 6-14, 7-14, 7-16

patchadd, 7-10

pdisk, 6-28, 6-31

pkginfo, 7-8

ps, 2-49, 3-43, 4-32, 5-55, 6-34, 7-45

pvcreate, 3-37

pvdisk, 3-29, 3-32, 3-37

raw, 5-42, 5-43, 5-51

rcp, 2-15, 3-15, 4-14, 7-16

reboot, 7-18

remsh, 3-16

rpm, 5-10

rsh, 2-15, 4-14, 7-17

sam, 3-18

scp, 2-15, 3-15, 4-14, 7-16

scstat, 7-45

service, 5-43, 5-51

setld, 4-7

setprivgrp, 3-16

shutdown, 4-17

sizer, 4-7

smit, 2-11, 2-12, 2-14

swap, 2-3, 3-3, 4-3, 5-3, 7-3

swapon, 2-3, 3-3, 4-3, 5-3, 7-3



- swlist, 3-7, 3-9
- sysconfig, 4-16
- sysconfigdb, 4-17
- sysctl, 5-19, 6-17
- sysdef, 7-18
- ulimit, 5-21
- umask, 2-50, 3-44, 4-33, 5-56, 6-35, 7-46
- uname, 3-7, 7-8
- unset, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- unsetenv, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- useradd, 3-13, 3-15, 4-12, 4-13, 5-14, 5-16, 7-14, 7-15
- varyoffvg, 2-45
- varyonvg, 2-42
- vgchange, 3-38, 3-40
- vgcreate, 3-38
- vgexport, 3-40
- vgimport, 3-40
- voldisk, 4-31
- vxassist, 7-42
- vxctl, 7-42
- vxvg, 7-40, 7-42
- vxdisk, 7-32, 7-40
- vxdiskadd, 7-40
- vxdiskconfig, 7-40
- vxvol, 7-42
- xhost, 2-2, 3-1, 4-1, 5-2, 7-2
- xterm, 2-2, 3-2, 4-2, 5-2, 7-2
- UNIX groups
  - checking for existing oinstall group, 2-10, 3-12, 4-10, 5-12, 6-7, 7-12
  - creating identical groups on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - creating the dba group, 2-11, 3-12, 4-10, 5-13, 6-8, 7-13
  - creating the oinstall group, 2-10, 3-11, 4-10, 5-12, 6-7, 7-12
  - creating the oper group, 2-11, 3-12, 4-11, 5-13, 6-9, 7-13
  - oinstall, 2-9, 3-11, 4-9, 5-11, 6-6, 7-11
  - OSDBA (dba), 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
  - OSOPER (oper), 2-9, 3-10, 4-9, 5-11, 6-6, 7-11
  - required for oracle user, 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - requirements, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
  - specifying when creating users, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
  - using NIS, 2-10, 2-13, 3-11, 3-14, 4-10, 4-12, 5-12, 5-15, 6-7, 6-12, 7-12, 7-15
- UNIX users
  - checking existence of the nobody user, 2-13, 4-12, 5-14, 6-11, 7-14
  - creating an unprivileged user on HP-UX, 3-14
  - creating identical users on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - creating the oracle user, 2-12, 3-13, 4-11, 5-13, 6-9, 7-13
  - Oracle Software Owner user (oracle), 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - setting shell limits for UNIX users on Linux, 5-20
  - setting up UNIX user equivalence, 2-15, 3-15, 4-14, 7-16, D-3
  - specifying UNIX groups when creating, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
  - UNIX nobody user, 4-9, 5-11, 6-6, 7-11
- requirements, 2-9, 3-10, 4-8, 5-11, 6-6, 7-11
- setting shell limits for on Linux, 5-20
- setting up user equivalence, 2-15, 3-15, 4-14, 7-16, D-3
- specifying groups when creating, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
- unprivileged user, 3-10, 4-9, 5-11, 6-6, 7-11
- using NIS, 2-10, 2-13, 3-11, 3-14, 4-10, 4-12, 5-12, 5-15, 6-7, 6-12, 7-12, 7-15
- UNIX workstation
  - installing from, 2-2, 3-1, 4-1, 5-2, 7-2
- unprivileged user
  - checking on AIX, Linux, Solaris, and Tru64 UNIX, 2-13, 4-12, 5-14, 7-14
  - creating on HP-UX, 3-14
  - nobody user, 6-6
- unset command, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- unsetenv command, 2-52, 3-46, 4-35, 5-58, 6-37, 7-48
- Upgrading an Existing Database page
  - Oracle Universal Installer, 11-7
- user accounts
  - post-installation setup, 13-5
- user equivalence
  - /etc/hosts.equiv file, 2-15, 3-15, 4-14, 7-16
  - .rhosts file, 2-15, 3-16, 4-14, 7-16
  - setting up, 2-15, 3-15, 4-14, 7-16
  - testing, 2-15, 3-16, 4-14, 7-17, 9-1, 11-3, 11-6
- user IDs
  - identifying existing, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
  - specifying, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
  - specifying on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
- USER tablespace
  - raw device for
    - on AIX, 2-35, 2-43
    - on HP-UX, 3-31, 3-39
    - on Linux, 5-45
    - on Solaris, 7-35, 7-41
- useradd command, 3-13, 3-15, 4-12, 4-13, 5-14, 5-16, 7-14, 7-15
- users
  - checking existence of the nobody user, 2-13, 4-12, 5-14, 6-11, 7-14
  - creating an unprivileged user on HP-UX, 3-14
  - creating identical users on other nodes, 2-13, 3-14, 4-12, 5-15, 6-12, 7-15
  - creating the oracle user, 2-12, 3-13, 4-11, 5-13, 6-9, 7-13
  - Oracle Software Owner user (oracle), 2-10, 3-11, 4-9, 5-12, 6-7, 7-11
  - setting shell limits for UNIX users on Linux, 5-20
  - setting up UNIX user equivalence, 2-15, 3-15, 4-14, 7-16, D-3
  - specifying UNIX groups when creating, 2-14, 3-14, 4-13, 5-15, 6-12, 7-15
  - UNIX nobody user, 4-9, 5-11, 6-6, 7-11
- USERS tablespace
  - described, 15-2
  - raw device, 8-19

## V

---

varyoffvg command, 2-45  
varyonvg command, 2-42  
Veritas Volume Manager, 7-32, 7-40  
vfs  
    kernel subsystem on Tru64 UNIX, 4-16  
vgchange command, 3-38, 3-40  
vgcreate command, 3-38  
vgexport command, 3-40  
vgimport command, 3-40  
VIP, 1-3, 8-7, 9-1, 11-11, 11-12  
VIPCA, 11-11  
    Configuration Results page, 11-11  
    IP Address page, 11-11  
    Public Network Interfaces page, 11-11  
    Summary page, 11-11  
    Welcome page, 11-11  
Virtual Internet Protocol Configuration Assistant  
    See VIPCA  
virtual IP  
    See VIP  
vm  
    kernel subsystem on Tru64 UNIX, 4-16  
voldisk command, 4-31  
volume group  
    creating a directory for on HP-UX, 3-38  
    creating on AIX, 2-41  
    creating on HP-UX, 3-37  
    creating on Solaris, 7-39  
    disabling activation on HP-UX, 3-36  
volume groups  
    creating on AIX, 2-42  
    creating on HP-UX, 3-38  
    creating on Solaris, 7-40  
volumes  
    creating logical volumes on AIX, 2-44  
    creating logical volumes on HP-UX, 3-39  
    creating logical volumes on Solaris, 7-42  
voting disk  
    raw device for, 6-30  
        on AIX, 2-35, 2-44  
        on HP-UX, 3-31, 3-39, 7-35  
        on Linux, 5-45  
        on Solaris, 7-42  
    restriction on using ASM for, 2-28, 3-24, 4-26,  
        5-31, 6-24, 7-28  
    specify location during install, 9-4  
    Windows, 10-2  
Voting Disk Information page  
    Oracle Universal Installer, 9-4  
vps\_ceiling parameter  
    recommended value on HP-UX, 3-17  
vxassist command, 7-42  
vxdctl command, 7-42  
vxdg command, 7-40, 7-42  
vxdisk command, 7-32, 7-40  
vxdiskadd command, 7-40  
vxdiskconfig command, 7-40  
VxVM, 7-32, 7-40  
    creating raw logical volumes on Solaris, 7-40

vxvol command, 7-42

## W

---

WebSphere MQ  
    checking requirements on Tru64 UNIX, 4-7  
    CSD download location  
        for AIX, 2-9  
        for HP-UX, 3-9  
        for Linux, 5-11  
        for Solaris, 7-11  
    CSDs required  
        on AIX, 2-8  
        on HP-UX, 3-9  
        on Linux, 5-9  
        on Solaris, 7-10  
        on Tru64 UNIX, 4-8  
    requirement on AIX, 2-6  
    requirement on HP-UX, 3-6  
    requirement on Linux, 5-9  
    requirement on Solaris, 7-7  
    requirement on Tru64 UNIX, 4-6  
Welcome page  
    Oracle Universal Installer, 9-3, 10-4, 11-5, 11-17  
    VIPCA, 11-11  
Windows  
    32-bit installation, 8-1  
    analogy for Oracle base directory, 2-18, 3-19, 4-17,  
        5-21, 6-18, 7-18  
    creating a raw device mapping file, 8-20  
    minimum memory installation, 11-12  
    system variables dialog, 11-13  
Windows Environment Variables page, 11-12, 11-13  
Windows ipconfig list, 8-8  
Windows raw partitions  
    create, 8-4  
Windows system variables dialog, 11-12  
Windows web browser support, 8-5  
wmem\_default file, 5-19  
wmem\_default parameter  
    recommended value on Linux, 5-19  
wmem\_max file, 5-19  
wmem\_max parameter  
    recommended value on Linux, 5-19

## X

---

X emulator  
    installing from, 2-2, 3-2, 4-2, 5-2, 7-2  
X library symbolic links  
    required on HP-UX, 3-10  
X Window and X/Motif Software development  
    Motif requirement, 4-6  
X window system  
    enabling remote hosts, 2-2, 3-1, 3-2, 4-1, 4-2, 5-2,  
        7-2  
X Windows Motif  
    requirement on Solaris, 7-7  
XFree86-devel  
    requirement on Linux, 5-9

xhost command, 2-2, 3-1, 4-1, 5-2, 7-2  
xterm command, 2-2, 3-2, 4-2, 5-2, 7-2

