

Oracle® Database

Installation Guide

10g Release 1 (10.1) for UNIX Systems: AIX-Based Systems, hp
HP-UX, hp Tru64 UNIX, Linux, and Solaris Operating System

Part No. B10811-05

October 2004

Oracle Database Installation Guide, 10g Release 1 (10.1) for UNIX Systems: AIX-Based Systems, hp HP-UX, hp Tru64 UNIX, Linux, and Solaris Operating System

Part No. B10811-05

Copyright © 1996, 2004, Oracle. All rights reserved.

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	ix
Preface	xi
Audience.....	xi
Documentation Accessibility	xi
Terminology	xii
Typographic Conventions.....	xii
Command Syntax	xii
Accessing Documentation.....	xiii
Related Documentation	xiv
Third Party Software Notices	xiv
1 Installation Overview	
Installation Overview	1-2
Oracle Database Installation Methods	1-2
Interactive Installation Methods	1-2
Automated Installation Methods Using Response Files	1-3
Oracle Database Installation Types	1-3
Database Configuration Options	1-4
New Installation Options in this Release	1-5
Database Storage Options	1-5
Database Management Options	1-6
Database Backup and Recovery Options.....	1-8
E-mail Notification Options.....	1-9
Installation Considerations	1-10
Hardware and Software Certification	1-10
Multiple Oracle Homes	1-10
Installing the Software on a System with an Existing Oracle Installation.....	1-10
Oracle Cluster Synchronization Services Daemon.....	1-10
Using Network Attached Storage or NFS File Systems	1-11
Upgrade Issues	1-11
Upgrading From Oracle8 Release 8.0.6 to Oracle Database 10g Release 1 (10.1).....	1-11
AL24UTFFSS Character Set	1-11
Upgrading an Existing Database on AIX.....	1-12

2 Pre-installation Tasks

Log In to the System as root	2-2
Check the Hardware Requirements	2-3
Check the Software Requirements	2-5
Checking the Software Requirements on AIX	2-6
Checking the Software Requirements on HP-UX.....	2-8
Checking the Software Requirements on Linux	2-14
Checking the Software Requirements on Solaris.....	2-19
Checking the Software Requirements on Tru64 UNIX	2-23
Create Required UNIX Groups and User	2-26
Creating the Oracle Inventory Group	2-27
Creating the OSDBA Group	2-28
Creating an OSOPER Group (Optional)	2-28
Creating the Oracle Software Owner User	2-29
HP-UX Only: Creating an Unprivileged User	2-31
Verifying that the UNIX User nobody Exists.....	2-31
HP-UX Only: Grant Privileges to the OSDBA Group.....	2-32
Configure Kernel Parameters and Shell Limits	2-33
Configuring Shell Limits and System Configuration Parameters on AIX.....	2-33
Configuring Kernel Parameters on HP-UX.....	2-34
Configuring Kernel Parameters on Linux	2-36
Configuring Kernel Parameters on Solaris.....	2-39
Configuring Kernel Subsystem Attributes on Tru64 UNIX.....	2-40
Identify Required Software Directories	2-41
Identify or Create an Oracle Base Directory	2-43
Choose a Storage Option for Oracle Database and Recovery Files	2-46
Create Directories for Oracle Database or Recovery Files	2-46
Configure Disks for Automatic Storage Management	2-49
Identifying Storage Requirements for ASM.....	2-49
Using an Existing ASM Disk Group.....	2-51
Configuring Disks for ASM.....	2-53
Configuring Disks for ASM on AIX	2-53
Configuring Disks for ASM on HP-UX.....	2-54
Configuring Disks for ASM on Linux	2-55
Configuring Disks for ASM Using the ASM Library Driver	2-55
Configuring Disks for ASM Using Raw Devices	2-60
Configuring Disks for ASM on Solaris.....	2-63
Configuring Disks for ASM on Tru64 UNIX.....	2-65
Configure Raw Devices	2-66
Configuring Raw Logical Volumes on AIX.....	2-67
Configuring Raw Logical Volumes on HP-UX.....	2-71
Configuring Raw Partitions or Raw Logical Volumes on Linux	2-75
Configuring Raw Partitions or Raw Logical Volumes on Solaris.....	2-84
Stop Existing Oracle Processes	2-92
Configure the oracle User's Environment	2-93

3	Installation Tasks	
	Installation Overview	3-2
	Download Oracle Software from the OTN Web Site	3-2
	Copy the Oracle Database Software to a Hard Disk	3-3
	AIX Only: Run the rootpre.sh Script	3-5
	Install the Oracle Database Software	3-5
	Reviewing Product-Specific Installation Guidelines	3-5
	Running the Oracle Universal Installer	3-6
4	Post-installation Tasks	
	Required Post-installation Tasks	4-2
	Downloading and Installing Patches	4-2
	Configuring Oracle Products.....	4-2
	HP-UX Only: Changing Permissions on the extjob Executable	4-2
	Recommended Post-installation Tasks	4-3
	Backing Up the root.sh Script.....	4-3
	Configuring New or Upgraded Databases.....	4-3
	Setting Up User Accounts	4-3
	Generating the Client Static Library.....	4-4
	Required Product-Specific Post-installation Tasks	4-4
	Configuring Oracle Net Services	4-4
	Configuring Oracle Label Security	4-5
	Installing Natively Compiled Java Libraries for Oracle JVM and Oracle <i>interMedia</i>	4-5
	Installing Oracle Text Supplied Knowledge Bases	4-6
	Configuring Oracle Messaging Gateway	4-6
	Modifying the listener.ora File for External Procedures	4-6
	Modifying the tnsnames.ora File for External Procedures	4-9
	Setting up the mgw.ora Initialization File.....	4-9
	Configuring Oracle Precompilers.....	4-10
	Configuring Pro*C/C++	4-10
	Linux x86 and Linux Itanium Only: Configure OCCI.....	4-10
	AIX, HP-UX, Solaris SPARC, and zSeries Linux: Configuring the 64-Bit Oracle JDBC/OCI Driver for JDK 1.4.....	4-11
	Installing Oracle Database 10g Products from the Companion CD.....	4-11
5	Reviewing Your Preconfigured Database	
	Log In to the Oracle Enterprise Manager Database Control.....	5-2
	Review Accounts and Passwords	5-3
	Unlock and Reset User Passwords	5-6
	Identify Databases	5-7
	Locate the Parameter File	5-8
	Review Tablespace and Datafiles, Redo Log Files, and Control Files	5-8
	Identifying Tablespaces and Datafiles	5-8
	Locating Redo Log Files	5-9
	Locating Control Files.....	5-10

6	Removing Oracle Software	
	Overview	6-2
	Identifying All Instances	6-2
	Removing an Oracle Database.....	6-2
	Removing an ASM Instance.....	6-3
	Reconfiguring Oracle Cluster Synchronization Services.....	6-4
	Removing Oracle Software.....	6-7
A	Mounting Discs	
	Mounting Discs on AIX	A-2
	Mounting Discs on HP-UX.....	A-2
	Mounting Discs on Linux.....	A-3
	Mounting Discs on Solaris.....	A-4
	Mounting Discs on Tru64 UNIX.....	A-5
B	Installing and Configuring Oracle Products Using Response Files	
	Introduction.....	B-2
	Creating a Database Using ASM as the Storage Option for Database Files	B-2
	Installation Overview	B-3
	Create the oraInst.loc File	B-3
	Prepare a Response File	B-4
	Run the Installer in Silent or Suppressed Mode.....	B-6
	Run NetCA in Silent Mode	B-7
	Run DBCA in Non-interactive or Silent Mode	B-8
C	Using NAS Devices	
	General Configuration Guidelines for NAS Devices	C-2
	Choosing Mount Points	C-2
	Choosing Mount Points for Oracle Software Files	C-2
	Choosing Mount Points for Oracle Database and Recovery Files	C-4
	Creating Files on a NAS Device for Use with ASM.....	C-5
	NFS Mount Options.....	C-6
D	Optimal Flexible Architecture	
	Changes to the Optimal Flexible Architecture for Oracle Database 10g.....	D-2
	Overview of the Optimal Flexible Architecture Standard.....	D-2
	Characteristics of an OFA-Compliant Installation.....	D-3
	Optimal Flexible Architecture Implemented on UNIX Systems	D-4
	File Systems.....	D-4
	Number of File Systems	D-4
	Naming Conventions	D-4
	Naming Mount Points for Very Large Databases (VLDBs).....	D-4
	Naming Directories.....	D-5
	Oracle Base Directory Naming Convention.....	D-5
	Referring to Pathnames.....	D-5

Oracle Home Directory Naming Convention.....	D-5
Naming Subdirectories	D-6
Naming Database Files.....	D-7
Separating Segments With Different Requirements	D-7
Naming Tablespaces.....	D-8
Exploiting the OFA Structure for Oracle Files	D-8
OFA File Mapping	D-9
E Default Port Numbers	
Components and Port Ranges.....	E-2
Changing the Enterprise Manager Agent HTTP Port	E-2
Changing the Enterprise Manager Database Control Ports.....	E-2
Changing the <i>iSQL*Plus</i> Ports	E-3
Changing the Ultra Search Ports	E-4
F Troubleshooting	
Verify Requirements.....	F-2
X Windows Display Errors	F-2
What to Do If an Installation Error Occurs	F-3
Reviewing the Log of an Installation Session	F-3
Troubleshooting Configuration Assistants	F-4
Configuration Assistant Failure.....	F-4
Fatal Errors.....	F-4
Silent-Mode Response File Error Handling.....	F-4
Cleaning Up After a Failed Installation.....	F-5

Index

Send Us Your Comments

**Oracle Database Installation Guide, 10g Release 1 (10.1) for UNIX Systems:
AIX-Based Systems, hp HP-UX, hp Tru64 UNIX, Linux, and Solaris Operating System**

Part No. B10811-05

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 errors or have other suggestions for improvement, you can e-mail them to us at osdwrite_us@oracle.com. Please indicate the title and part number of the document and the chapter, section, and page number if available. Please indicate whether you would like a reply.

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

Preface

This guide describes how to install and configure Oracle Database 10g release 1 (10.1) on UNIX systems.

Audience

This guide is intended for anyone responsible for installing Oracle Database 10g release 1 (10.1) on a single UNIX system. If you want to install Oracle Real Application Clusters on a cluster, see the *Oracle Real Application Clusters Installation and Configuration Guide* for installation instructions.

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.

Terminology

The names for UNIX operating systems have been shortened in this guide, as follows:

Operating System	Abbreviated Name
AIX-Based Systems	AIX
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.
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.

Typographic Conventions

The following typographic conventions are used in this guide:

Convention	Description
<code>monospace</code>	Monospace type indicates UNIX commands, directory names, usernames, pathnames, and filenames.
<i>italics</i>	Italic type indicates variables, including variable portions of filenames. It is also used for emphasis and for book titles.
UPPERCASE	Uppercase letters indicate Structured Query Language (SQL) reserved words, initialization parameters, and environment variables.

Command Syntax

UNIX command syntax appears in `monospace` font. The dollar character (\$), number sign (#), or percent character (%) are UNIX command prompts. Do not enter them as part of the command. The following command syntax conventions are used in this guide:

Convention	Description
backslash \	A backslash is the UNIX command continuation character. It is used in command examples that are too long to fit on a single line. Enter the command as displayed (with a backslash) or enter it on a single line without a backslash: <pre>dd if=/dev/rdisk/c0t1d0s6 of=/dev/rst0 bs=10b \ count=10000</pre>
braces { }	Braces indicate required items: <pre>.DEFINE {macro1}</pre>
brackets []	Brackets indicate optional items: <pre>cvtcrt termname [outfile]</pre>

Convention	Description
ellipses ...	Ellipses indicate an arbitrary number of similar items: CHKVAL fieldname value1 value2 ... valueN
<i>italics</i>	Italic type indicates a variable. Substitute a value for the variable: <i>library_name</i>
vertical line	A vertical line indicates a choice within braces or brackets: FILE <i>filesize</i> [K M]

Accessing Documentation

The documentation for Oracle Database 10g release 1 (10.1) for UNIX systems includes platform-specific documentation and generic product documentation.

Platform-Specific Documentation

Platform-specific documentation includes information about installing and using Oracle products on particular platforms. The platform-specific documentation for this product is available in both Adobe portable document format (PDF) and HTML format on the product disc. To access the platform-specific documentation on disc:

1. Use a Web browser to open the `welcome.htm` file in the top-level directory of the disc.
2. For DVD-ROMs only, select the appropriate product link.
3. Select the **Documentation** tab.

If you prefer paper documentation, then open and print the PDF files.

Product Documentation

Product documentation includes information about configuring, using, or administering Oracle products on any platform. The product documentation for Oracle Database 10g products is available in both HTML and PDF formats in the following locations:

- On the Oracle Database 10g Documentation Library CD-ROM
To access the documentation from the CD-ROM, use a Web browser to view the `index.htm` file in the top-level directory on the disc.
- In the `doc` subdirectory on the Oracle Database 10g DVD-ROM
To access the documentation from the DVD-ROM, use a Web browser to view the `welcome.htm` file in the top-level directory on the disc, then select the Oracle Database 10g Documentation Library link.
- Online on the Oracle Technology Network (OTN) Web site:
<http://www.oracle.com/technology/documentation>

Related Documentation

The platform-specific documentation for Oracle Database 10g products includes the following manuals:

- Oracle Database:
 - *Oracle Database Release Notes (platform-specific)*
 - *Oracle Database Quick Installation Guide (platform-specific)*
 - *Oracle Database Installation Guide for UNIX Systems*
 - *Oracle Real Application Clusters Installation and Configuration Guide*
 - *Oracle Database Administrator's Reference for UNIX Systems*
 - *Oracle Procedural Gateway for APPC Installation and Configuration Guide for UNIX*
 - *Oracle Procedural Gateway for APPC User's Guide for UNIX*
 - *Oracle Procedural Gateway for APPC Messages Guide*
 - *Oracle Transparent Gateway for DRDA Installation and User's Guide for UNIX*
- Oracle Client:
 - *Oracle Database Client Quick Installation Guide (platform-specific)*
 - *Oracle Database Client Installation Guide for UNIX Systems*
- Oracle Database 10g Companion CD:
 - *Oracle Database Companion CD Installation Guide for UNIX Systems*
 - *Oracle Database Companion CD Quick Installation Guide (platform-specific)*

Refer to the Oracle Database release notes for your platform for important information that was not available when this book was released. The release notes for Oracle Database 10g are updated regularly. You can get the most-recent version from OTN:

<http://www.oracle.com/technology/documentation>

Third Party Software Notices

This program contains third party software from HP. The Oracle program license that accompanied this product determines your right to use the Oracle program, including the HP software. Notwithstanding anything to the contrary in the Oracle program license, the HP software is provided "AS IS" and without intellectual property indemnities, warranties, or support of any kind from Oracle or HP.

This program contains third party software from International Business Machines Corporation ("IBM"). The Oracle program license that accompanied this product determines your right to use the Oracle program, including the IBM software.

Notwithstanding anything to the contrary in the Oracle program license, the IBM software is provided "AS IS" and without intellectual property indemnities, warranties, or support of any kind from Oracle or IBM.

Installation Overview

This chapter describes the different types of Oracle Database installations that you can perform, and issues that you should consider before installing the software. It includes information about the following topics:

- [Installation Overview](#)
- [Oracle Database Installation Methods](#)
- [Oracle Database Installation Types](#)
- [New Installation Options in this Release](#)
- [Installation Considerations](#)
- [Upgrade Issues](#)

Installation Overview

The Oracle Database installation process consists of four parts:

1. **Planning your installation:** This chapter describes the Oracle products that you can install and issues that you must consider before starting the installation.
2. **Completing pre-installation tasks:** [Chapter 2](#) describes pre-installation tasks that you must complete before installing the product.
3. **Installing software:** [Chapter 3](#) describes how to use the Oracle Universal Installer to install this product.
4. **Completing post-installation tasks:** [Chapter 4](#) describes recommended and required post-installation tasks.

Oracle Database Installation Methods

You can choose different installation methods to install Oracle Database, as follows:

- [Interactive Installation Methods](#)
- [Automated Installation Methods Using Response Files](#)

Interactive Installation Methods

When you use the interactive method to install Oracle Database, the Installer displays a series of screens that enable you to specify all of the required information to install the Oracle Database software and optionally create a database.

With Oracle Database 10g release 10.1.0.3 on Linux x86 or Linux Itanium, the Installer provides two interactive methods that you can use to install Oracle Database:

- **Basic Installation method**

Select this installation method if you want to quickly install Oracle Database. This installation method requires minimal user input. It installs the software and optionally creates a general-purpose database using the information that you specify on one screen. It is the default installation method.
- **Advanced Installation method**

Select this installation method if you want to complete any of the following tasks:

 - Perform a custom software installation, or choose a different database configuration
 - Install Oracle Real Application Clusters
 - Upgrade an existing database
 - Select a database character set or different product languages
 - Create the EXAMPLE tablespace during the installation
 - Create a database on a different file system from the software
 - Configure Automatic Storage Management (ASM) or use raw devices for database storage
 - Specify different passwords for administrative schemas
 - Configure automated backups or Oracle Enterprise Manager notifications

Note: To choose the Custom installation type or Advanced database configuration options mentioned in this guide, you must choose the Advanced Installation method.

With release 10.1.0.2 on all platforms or release 10.1.0.3 on platforms other than Linux x86 and Linux Itanium, the Basic Installation method is not available. The Advanced Installation method is used by default.

Automated Installation Methods Using Response Files

By creating a response file and specifying this file when you start the Installer, you can automate some or all of the Oracle Database installation. These automated installation methods are useful if you need to perform multiple installations on similarly configured systems or if the system where you want to install the software does not have X Window system software installed.

When you use a response file, you can run the Installer in the following modes, depending on whether you specify all of the required information or not:

- **Silent Mode:** The Installer runs in silent mode if you use a response file that specifies all required information. None of the Installer screens are displayed.
- **Suppressed Mode:** The Installer runs in suppressed mode if you do not specify all required information in the response file. The Installer displays only the screens that prompt for the information that you did not specify.

For more detailed information about these modes and about how to complete an installation using response files, see [Appendix B](#).

Oracle Database Installation Types

Note: Oracle Client is installed separately. You cannot install Oracle Client during an Oracle Database installation.

You can choose one of the following installation types when installing Oracle Database:

Note: For more information about the features available with each Oracle Database edition and for information about licensing, see the *Oracle Database Licensing Information* guide.

- **Standard Edition:** Installs an integrated set of management tools, full distribution, replication, Web features, and facilities for building business-critical applications.
- **Enterprise Edition:** Installs licensable Oracle Database options and database configuration and management tools in addition to all of the products that are installed during a Standard Edition installation. It also installs products most commonly used for data warehousing and transaction processing.
- **Custom:** Enables you to select the individual components that you want to install from the list of all available components.

Database Configuration Options

During the installation, you can choose whether you want to create an Oracle database as part of the installation. If you choose to create an Oracle database, the Installer uses the Database Configuration Assistant (DBCA) to create it. You can choose to create one of the preconfigured database types, which are designed for a variety of different applications, modify one of the preconfigured database types, or create a customized database to suit your own requirements.

Preconfigured Database Types

Oracle provides the following preconfigured database types that you can create or customize during the installation:

- General Purpose
- Transaction Processing
- Data Warehouse

See the online help provided by either the Installer or the DBCA for a description of these preconfigured database types.

Installation Choices that Affect Database Creation

The Installer runs DBCA in one of two modes, depending on the choices that you make during the installation:

- Non-interactive mode

If you choose either the Enterprise Edition or Standard Edition installation type, then choose to create a preconfigured database type, the Installer prompts you for the minimum amount of information required to create a database of the type you choose. It then runs DBCA in non-interactive mode to create the database after it installs the software.

Note: Oracle recommends that you use this method to create a database if you have not previously created one.

- Interactive mode

If you choose the Custom installation type or choose the Advanced database configuration option, the Installer does not prompt you for database information. Instead, it installs the software and then runs DBCA in interactive mode. Using the screens in DBCA, you can either modify one of the preconfigured database types or create a custom database and specify precisely how you want to configure it.

Note: If you choose this method to create a database, click **Help** on any of the DBCA screens for a description of the information that you must specify on that screen.

Creating a Database After Installation

If you decide not to create a database during the installation, you can use DBCA to create one after you have installed the software. For more information about using DBCA to create a database after installation, see the *Oracle Database 2 Day DBA* manual.

New Installation Options in this Release

This release of Oracle Database provides new options that you can choose during installation to simplify database administration tasks. These options include:

- [Database Storage Options](#)
- [Database Management Options](#)
- [Database Backup and Recovery Options](#)
- [E-mail Notification Options](#)

The following sections describe these options.

Database Storage Options

If you choose to create a database during the installation, you can specify one of three storage options for database files:

- [File System](#)
- [Automatic Storage Management](#)
- [Raw Devices](#)

File System

If you choose the file system option, the Database Configuration Assistant creates the database files in a directory on a file system mounted on your computer. Oracle recommends that the file system you choose be separate from the file systems used by the operating system or the Oracle software. The file system that you choose can be any of the following:

- A file system on a disk that is physically attached to the system
If you are creating a database on basic disks that are not logical volumes or RAID devices, Oracle recommends that you follow the Optimal Flexible Architecture (OFA) recommendations described in [Appendix D](#) and distribute the database files over more than one disk.
- A file system on a logical volume manager (LVM) volume or a RAID device
If you are using multiple disks in an LVM or RAID configuration, Oracle recommends that you use the stripe and mirror everything (SAME) methodology to increase performance and reliability. Using this methodology, you do not need to specify more than one file system mount point for database storage.
- A network file system (NFS) mounted from a certified network attached storage (NAS) device

You can store database files on NAS devices provided that the NAS device is certified by Oracle. See the "[Using Network Attached Storage or NFS File Systems](#)" section on page 1-11 for more information about certified NAS and NFS devices.

If you choose the Custom installation type or the Advanced database creation option, you can also choose to use the Oracle-managed files feature with the new database. If you use this feature, you need only specify the database object name instead of file names when creating or deleting database files.

See Also: For more information about Oracle-managed files, see the *Oracle Database Administrator's Guide*.

Automatic Storage Management

Automatic Storage Management (ASM) is a new feature introduced with this release of Oracle Database. It is a high-performance storage management solution for Oracle database files that is consistent across all supported platforms. Designed specifically to simplify the job of the database administrator (DBA), ASM provides you with a flexible storage solution that simplifies the management of a dynamic database environment. The features provided by ASM make most manual I/O performance tuning tasks unnecessary.

To use ASM for database storage, you must create one or more ASM disk groups. A disk group is a set of disk devices that ASM manages as a single unit. ASM spreads data evenly across all of the devices in the disk group to optimize performance and utilization. To protect against disk failure, you can choose one of three redundancy levels when you create a disk group. The redundancy level defines how files are mirrored within a disk group, as follows:

Redundancy Level	Mirroring
Normal	Two-way mirroring
High	Three-way mirroring
External	No mirroring by ASM

In addition to the manageability, performance, and reliability benefits provided by ASM, it can also increase database availability. You can add or remove disk devices from disk groups without shutting down the database. ASM automatically rebalances the files across the disk group after disks have been added or removed.

Disk groups are managed by a special Oracle instance, called an ASM instance. This instance must be running before you can start a database instance that uses ASM for storage management. If you choose ASM as the storage mechanism for your database, DBCA creates and starts this instance if necessary.

See Also: For a more detailed description of ASM, see the *Oracle Database Concepts* manual. For information about administering ASM, see the *Oracle Database Administrator's Guide*.

Raw Devices

Raw devices are disk partitions or logical volumes that have not been formatted with a file system. When you use raw devices for database file storage, Oracle writes data directly to the partition or volume, bypassing the operating system file system layer. For this reason, you can sometimes achieve performance gains by using raw devices. However, because raw devices can be difficult to create and administer, and because the performance gains over modern file systems are minimal, Oracle recommends that you choose ASM or file system storage in preference to raw devices.

Database Management Options

To simplify database administration, Oracle provides a Web-based management tool called Oracle Enterprise Manager. There are two ways that you can deploy Oracle Enterprise Manager, as follows:

- Deploy Oracle Enterprise Manager 10g centrally in your environment
 - To deploy Oracle Enterprise Manager centrally, you must install at least one Oracle Management Repository and one Oracle Management Service within your environment, then install an Oracle Management Agent on every computer that

you want to manage. You can then use a single Web-based interface to manage and monitor software and hardware targets on all of those systems. Targets can include Oracle databases, application servers, Net listeners, and third-party software. This single interface is called Oracle Enterprise Manager Grid Control (or simply Grid Control).

Note: Oracle Enterprise Manager 10g is available separately in the Oracle Database media pack. For more information about Oracle Enterprise Manager 10g, see the *Oracle Enterprise Manager 10g Concepts* manual and the *Oracle Enterprise Manager 10g Installation and Basic Configuration* manual.

- Deploy Oracle Enterprise Manager locally on the database system

Oracle Enterprise Manager software is installed by default¹ with every Oracle Database installation. This local installation provides a Web-based interface called Oracle Enterprise Manager Database Control. The Database Control is similar in function to the Grid Control, but it can manage only a single database. If you want to administer more than one database on this system, you must either configure a separate Database Control for each one, or install Oracle Enterprise Manager 10g.

Management Options for Preconfigured Databases

When you choose to create a preconfigured database during the installation, you must select the Oracle Enterprise Manager interface that you want to use to manage the database. The following options are available:

- Use Grid Control for database management

This option is available only if an Oracle Management Agent is installed on the system. When the Installer detects an Oracle Management Agent on the system, it allows you to choose this option and specify the Oracle Management Service that you want to use to manage the database.

If an Oracle Management Agent is not installed, you must choose to use Database Control to manage the database. However, if you install Oracle Management Agent after you install Oracle Database, you can then use Grid Control to manage this database.

- Use Database Control for database management

This option is selected by default if an Oracle Management Agent is not installed on the system. However, even if a Management Agent is installed, you can still choose to configure Database Control to manage the database.

Management Options for Custom Databases

If you choose the Advanced database configuration option or choose to create a database during a Custom installation, the Installer runs Database Configuration Assistant (DBCA) in interactive mode. Using a screen in DBCA, you can specify the Oracle Enterprise Manager interface that you want to use to manage the database. Alternatively, you can also choose not to configure the database with Enterprise Manager.

¹ During a custom installation, you can choose not to install Oracle Enterprise Manager software, however, Oracle recommends that you do install it.

Note: Oracle recommends that you configure the database to use Enterprise Manager during installation. However, if you choose not to configure the database to use Enterprise Manager during the installation, you can use DBCA after the installation to configure the database to use it.

Features Provided by Oracle Enterprise Manager Database Control

Oracle Enterprise Manager Database Control provides a Web-based user interface that enables you to monitor, administer, and maintain an Oracle database. You can use it to perform all of your database administration tasks. You can also use it to determine information about the database, such as:

- Instance name, database version, Oracle home location, media recovery options, and other instance data
- Current instance availability
- Database alert information
- Session and SQL-related performance information
- Space usage metrics

Furthermore, it provides you with automatic notification of security alerts and it provides the ability to download and apply patches for the software.

Database Backup and Recovery Options

If you choose to use Oracle Enterprise Manager Database Control during the installation, you can optionally enable automated database backups that use the Oracle-suggested default backup strategy.

Note: You do not have to enable automated backups during the installation. If you prefer, you can use Oracle Enterprise Manager Database Control or Grid Control to configure automated backups after you install the software and create a database.

Enabling Automated Backups

If you enable automated backups, Oracle Enterprise Manager schedules a daily backup job that uses Oracle Recovery Manager (RMAN) to back up all of the database files to an on-disk storage area called the flash recovery area. The first time the backup job runs, it creates a full backup of the database. Subsequent backup jobs perform incremental back-ups, which enable you to recover the database to its state at any point during the preceding 24 hours.

To enable automated backup jobs during installation, you must specify the following information:

- The location of the flash recovery area

You can choose to use either a file system directory or an ASM disk group for the flash recovery area. The default disk quota configured for the flash recovery area is 2 GB. For ASM disk groups, the required disk space depends on the redundancy level of the disk group that you choose. [Chapter 2](#) describes how to choose the location of the flash recovery area and identifies its disk space requirements.

- An operating system user name and password for the backup job

Oracle Enterprise Manager uses the operating system credentials that you specify when running the backup job. The user name that you specify must belong to the UNIX group that identifies database administrators (the OSDBA group, typically `dba`). The Oracle software owner user (typically `oracle`) that you use to install the software is a suitable choice for this user. [Chapter 2](#) describes the requirements for the OSDBA group and Oracle software owner user and describes how to create them.

Backup Job Default Settings

If you enable automated backups after choosing one of the preconfigured databases during the installation, automated backup is configured with the following default settings:

- The backup job is scheduled to run nightly at 2 a.m.
- The disk quota for the flash recovery area is 2 GB.

If you enable automated backups by using DBCA, either during or after the installation, you can specify a different start time for the backup job and a different disk quota for the flash recovery area.

More Information About Automated Backups and Recovery

For information about using Oracle Enterprise Manager Database Control to configure or customize automated backups or to recover a backed up database, see the *Oracle Database 2 Day DBA* manual.

For more detailed information about defining a back-up strategy and backing up and recovering Oracle databases, see the *Oracle Database Backup and Recovery Basics* manual or the *Oracle Database Backup and Recovery Advanced User's Guide*.

E-mail Notification Options

If you choose to use the Oracle Enterprise Manager Database Control during the installation, you can configure Enterprise Manager to send e-mail when specific events occur. These events can include occurrences such as disk space reaching a critical limit (a threshold), or a database shutting down unexpectedly.

If you choose to enable e-mail notifications, you must specify the following information:

- The host name of an simple mail transport protocol (SMTP) server
- The e-mail address that should receive the alerts

The e-mail address that you specify could belong to an individual or it could be a shared e-mail account or a distribution list.

You can use Enterprise Manager Database Control to set-up, change, or customize e-mail notifications after you have created the database.

Installation Considerations

This section contains information that you should consider before deciding how to install this product. It contains the following sections:

- [Hardware and Software Certification](#)
- [Multiple Oracle Homes](#)
- [Oracle Cluster Synchronization Services Daemon](#)
- [Using Network Attached Storage or NFS File Systems](#)

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 *OracleMetaLink* Web site for the most up-to-date list of certified hardware platforms and operating system versions. The *OracleMetaLink* Web site is available at the following URL:

<http://metalink.oracle.com>

If you do not have a current Oracle Support Services contract, you can access the same information from the following Web site:

<http://www.oracle.com/technology/support/metalink/content.html>

Multiple Oracle Homes

This product supports multiple Oracle homes. This means that you can install this release or previous releases of the software more than once on the same system, in different Oracle home directories.

Installing the Software on a System with an Existing Oracle Installation

You must install this product into a new Oracle home directory. You cannot install products from one release of Oracle Database into an Oracle home directory of a different release. For example, you cannot install release 10.1 software into an existing Oracle9i Oracle home directory. If you attempt to install this release into an Oracle home directory that contains software from an earlier Oracle release, the installation fails.

You can install this release more than once on the same system provided that each installation is installed in a separate Oracle home directory.

Oracle Cluster Synchronization Services Daemon

The first time you install Oracle Database 10g on a system, the installation configures and starts a single-node version of the Oracle Cluster Synchronization Services (CSS) daemon. The CSS daemon is required to enable synchronization between an ASM instance and the database instances that rely on it for database file storage. It is configured and started even if you do not choose ASM as a storage mechanism for database files. Because it must be running before any ASM instance or database instance starts, it is configured to start automatically when the system boots.

For Oracle Real Application Clusters installations, the CSS daemon is installed with Oracle Cluster Ready Services (CRS) in a separate Oracle home directory (also called the CRS home directory). For single-node installations, the CSS daemon is installed in

and runs from the same Oracle home as Oracle Database. For this reason, you must use caution when removing Oracle Database 10g software from the system. Before you remove an Oracle home directory that contains Oracle Database 10g, you must either delete the CSS daemon configuration, or if necessary, reconfigure the CSS daemon to run from another Oracle home directory.

See Also: For more information about deleting or reconfiguring the Oracle CSS daemon, see the "[Reconfiguring Oracle Cluster Synchronization Services](#)" section on page 6-4.

Note: If you plan to have more than one Oracle Database 10g installation on a single system and you want to use ASM for database file storage, Oracle recommends that you run the CSS daemon and the ASM instance from the same Oracle home directory and use different Oracle home directories for the database instances.

Using Network Attached Storage or NFS File Systems

Oracle Database must be able to verify that writes to a disk are completed successfully. NFS file systems, including file systems on NAS devices, may not be able to guarantee that writes to a disk are completed successfully, and this may lead to possible datafile corruption. Oracle recommends that you do not store files on NFS mounted file systems unless the storage vendor and storage device are listed in the Oracle Storage Compatibility Program list. This list is available from the following Web site:

<http://www.oracle.com/technology/deploy/availability/htdocs/oscp.html>

If your storage device is supported, you can use it to store Oracle software files, Oracle database files, or both.

See Also: For guidelines about using NFS and NAS devices for Oracle software or database files, see [Appendix C](#).

Upgrade Issues

For information about upgrading a previous release of Oracle Database to Oracle Database 10g, see the *Oracle Database Upgrade Guide*. The following sections include additional platform-specific upgrade information that you should review before upgrading an existing database.

Upgrading From Oracle8 Release 8.0.6 to Oracle Database 10g Release 1 (10.1)

If you are upgrading a database from Oracle8 release 8.0.6 to Oracle Database 10g and ConText Cartridge is installed, then you cannot use the Database Upgrade Assistant to upgrade the database. You must upgrade the database manually. For information about manually upgrading ConText Cartridge to Oracle Text, see the *Oracle8i interMedia Text Migration* manual and the *Oracle Database Upgrade Guide*.

AL24UTFFSS Character Set

Before you upgrade an existing database that uses the AL24UTFFSS character set, you must first upgrade the database character set to UTF8. Oracle recommends that you use the Character Set Scanner (`csscan`) utility for data analysis before attempting to upgrade the existing database character set.

The Character Set Scanner utility checks all character data in the database and tests for the effects of, and problems with, changing the character set encoding. Before running the Character Set Scanner utility, set the shared library path environment variable for your platform to include the `$ORACLE_HOME/lib` directory. The following table lists the appropriate shared library path environment path for each platform:

Platform	Environment Variable
AIX	LIBPATH
HP-UX, Linux, Solaris, and Tru64 UNIX	LD_LIBRARY_PATH

Upgrading an Existing Database on AIX

The following table describes the steps that you must follow to upgrade particular Oracle Database releases running on different versions of AIX. In some cases, to complete the upgrade, you must upgrade to an intermediate Oracle Database release before upgrading to Oracle Database 10g.

Oracle Database Release	AIX Version	Upgrade Instructions
8.0.6 or earlier	4.3.3	<ol style="list-style-type: none"> 1. Upgrade to Oracle8i release 8.1.7 (32-bit) for AIX by following the instructions in the <i>Oracle8i Migration</i> guide. 2. Upgrade the operating system to AIX 5L version 5.2. 3. Upgrade to Oracle Database 10g by following the instructions in the <i>Oracle Database Upgrade Guide</i>.
8.1.7 (64-bit) or 9.0.1	4.3.3	<ol style="list-style-type: none"> 1. Upgrade to Oracle9i release 2 (9.2) on AIX version 4.3.3 by following the instructions in the <i>Oracle9i Database Migration</i> guide. 2. Follow the instructions in the following row.
9.2	4.3.3	<ol style="list-style-type: none"> 1. Upgrade to Oracle9i release 2 (9.2) on AIX 5L version 5.2 by following the instructions in the <i>Oracle9i Release Notes for AIX-Based 5L Systems (64-bit)</i> and the <i>Oracle9i Release 2 Installation Guide for UNIX Systems</i>. 2. Upgrade to Oracle Database 10g by following the instructions in the <i>Oracle Database Upgrade Guide</i>.
8.1.7 (32-bit)	4.3.3, 5.1, or 5.2	<ol style="list-style-type: none"> 1. If necessary, upgrade the operating system to AIX 5L version 5.2. 2. Upgrade to Oracle Database 10g by following the instructions in the <i>Oracle Database Upgrade Guide</i>.
9.2	5.1 or 5.2	<ol style="list-style-type: none"> 1. If necessary, upgrade the operating system to AIX 5L version 5.2. 2. Upgrade to Oracle Database 10g by following the instructions in the <i>Oracle Database Upgrade Guide</i>.

Pre-installation Tasks

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 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](#)
- [Choose a Storage Option for Oracle Database and Recovery Files](#)
- [Create Directories for Oracle Database or Recovery Files](#)
- [Configure Disks for Automatic Storage Management](#)
- [Configure Raw Devices](#)
- [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, see [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` 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, see your X server documentation for more information about completing this procedure. Depending on the X server software that you are using, you may need to complete the tasks in a different order.

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

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

Check the Hardware Requirements

The system must meet the following minimum hardware requirements:

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

On systems with 2 GB or more of RAM, the swap space can be between one and two times the size of RAM. On AIX systems with 1 GB or more of memory, do not increase the swap space above 2 GB. See the *Oracle Database Administrator's Reference for UNIX Systems* for more information about configuring swap space on AIX.
- 400 MB of disk space in the `/tmp` directory
- Between 1.5 GB and 3 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 the system meets these requirements, follow these steps:

1. To determine the physical RAM size, enter one of the following commands:

Platform	Command
AIX	<code># /usr/sbin/lssattr -E -l sys0 -a realmem</code>
HP-UX	<code># grep "Physical:" /var/adm/syslog/syslog.log</code>
Linux	<code># grep MemTotal /proc/meminfo</code>
Solaris	<code># /usr/sbin/prtconf grep "Memory size"</code>
Tru64 UNIX	<code># /bin/vmstat -P grep "Total Physical Memory"</code>

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 one of the following commands:

Platform	Command
AIX	<code># /usr/sbin/lssps -a</code>
HP-UX	<code># /usr/sbin/swapinfo -a</code>
Linux	<code># grep SwapTotal /proc/meminfo</code>
Solaris	<code># /usr/sbin/swap -s</code>
Tru64 UNIX	<code># /sbin/swapon -s</code>

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

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

- HP-UX:

```
# bdf /tmp
```

- Other operating systems:

```
# 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 one of the following commands:

- HP-UX:

```
# bdf
```

- Other operating systems:

```
# df -k
```

The following table shows the approximate disk space requirements for software files for each installation type:

Platform	Installation Type	Requirement for Software Files (GB)
AIX	Enterprise Edition	3
	Standard Edition	3
	Custom (maximum)	3
HP-UX	Enterprise Edition	2.5
	Standard Edition	2.5
	Custom (maximum)	2.5
Linux	Enterprise Edition	1.5
	Standard Edition	1.5
	Custom (maximum)	1.5
Solaris	Enterprise Edition	1.5
	Standard Edition	1.5
	Custom (maximum)	1.5

Platform	Installation Type	Requirement for Software Files (GB)
Tru64 UNIX	Enterprise Edition	3
	Standard Edition	3
	Custom (maximum)	3

5. To determine whether the system architecture can run the software, enter one of the following commands:

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

Platform	Command	Expected Output
AIX	# /usr/bin/getconf HARDWARE_BITMODE	64
HP-UX	# /bin/getconf KERNEL_BITS	64
Linux (x86 and Itanium)	# 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.
Solaris	# /bin/isainfo -kv	SPARC systems: 64-bit sparcv9 kernel modules x86 systems: 32-bit i386 kernel modules

Check the Software Requirements

Depending on your operating system, see one of the following sections for information about checking the software requirements:

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

- [Checking the Software Requirements on AIX](#) on page 2-6
- [Checking the Software Requirements on HP-UX](#) on page 2-8
- [Checking the Software Requirements on Linux](#) on page 2-14
- [Checking the Software Requirements on Solaris](#) on page 2-19
- [Checking the Software Requirements on Tru64 UNIX](#) on page 2-23

Checking the Software Requirements on AIX

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> </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>
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:</p> <p><code>vac.C</code></p> <p>VisualAge C++ Compiler V6.0:</p> <p><code>vacpp.cmp.core</code></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
```

If a fileset is not installed and committed, then install it. See 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>

Installation Type or Product	Requirement
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>

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:

```
# lspp -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, see the following Web site for download and installation information:

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

4. To continue completing pre-installation tasks, go to the "[Create Required UNIX Groups and User](#)" section on page 2-26.

Checking the Software Requirements on HP-UX

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.
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
Oracle JDBC/OCI Drivers	You can use the following optional Java SDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation: <ul style="list-style-type: none"> ■ HP-UX 11i (11.11): <ul style="list-style-type: none"> ■ Java SDK 1.3.1.02 with the JNDI extension ■ Java SDK 1.2.2.09 ■ HP-UX 11i v2 (11.23): <ul style="list-style-type: none"> ■ Java SDK 1.3.1.09 Note: Java SDK 1.4.2.00 (PA-RISC) or 1.4.1.03 (Itanium) is installed with this release.
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE (PA-RISC only)	DCE (B.11.11.10): DCE-C-Tools DCE-CDS-Server DCE-Core DCE-CoreAdmin DCE-CoreTools DCE-SEC-Server

Installation Type or Product	Requirement
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. See your operating system or software documentation for information about installing products.

Check for Required Patches

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

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

Installation Type or Product	Requirement
All installations	Quality Pack bundle: HP-UX 11i (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, LDCCD access type <p>Note: See the following Web site for information about additional patches that might be required by JDK 1.4.2: http://www.hp.com/products1/unix/java/patches/index.html</p>

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 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, see the following Web site for download and installation information:

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

Create Required Symbolic Links

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

To allow 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 libXhp1.3 libXhp1.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
```

To continue completing pre-installation tasks, go to the ["Create Required UNIX Groups and User"](#) section on page 2-26.

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 2–1, "Software Requirements for Installations on x86 Systems"](#)
- [Table 2–2, "Software Requirements for Installations on Itanium Systems"](#)
- [Table 2–3, "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 2–4, "Product-Specific Software Requirements for Linux Systems"](#) is installed.

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

Table 2–1 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: 2.4.21-138</p> <p>SUSE Linux Enterprise Server 9: 2.6.5-7.5</p>

Table 2–1 (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 gcc-2.96-128 gcc-c++-2.96-128 libstdc++-2.96-128 openmotif-2.1.30-11</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 2–2 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 2–2 (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 2–3 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 2–4 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 2–4 (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>

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

<http://www.ibm.com/software/integration/mqfamily/support/summary/lin.html>

5. To continue completing pre-installation tasks, go to the "[Create Required UNIX Groups and User](#)" section on page 2-26.

Checking the Software Requirements on Solaris

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 SUNWlibm SUNWlibms SUNWsprot SUNWtoo SUNWilof SUNWilcs SUNWil5cs SUNWxwfont</p> <p>SPARC systems only:</p> <p>SUNWsprox</p> <p>Note: You might also require additional font packages for Java, depending on your locale. See the following Web site for more information:</p> <p>http://java.sun.com/j2se/1.4.2/font-requirements.html</p>
Oracle Spatial	<p>Motif RunTime Kit:</p> <p>SUNWmfrun</p> <p>Note: This software is required only to build the sample programs.</p>
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)

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)	Sun ONE Studio 8 (C and C++ 5.5)
Oracle JDBC/OCI Drivers	<p>You can use the following optional JDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation:</p> <ul style="list-style-type: none"> ■ Sun Java 2 SDK Standard Edition 1.3.1_10 (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, see 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 SUNWlibm SUNWlibms SUNWsprot \
SUNWsprox SUNWtoo SUNWilof SUNWilcs SUNWil5cs SUNWxwft
```

If a package is not installed, then install it. See your operating system or software documentation for information about installing packages.

Check for Required Patches

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

Note: The patch versions shown in the following table are minimum versions. Higher versions of the same patches are also supported.

Installation Type or Product	Requirement
All installations	<p>Patches for Solaris 8 (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 patch¹ ■ 108773-18, SunOS 5.8: IIIM and X I/O Method patch¹ ■ 108921-16, CDE 1.4: dtwm patch¹ ■ 108940-53, Motif 1.2.7 and 2.1.1: Runtime lib. patch for Solaris 8¹ ■ 108987-13, SunOS 5.8: Patch for patchadd and patchrm ■ 108989-02, /usr/kernel/sys/acctctl & /.../exacctsyst patch ■ 108993-18, 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
Oracle Messaging Gateway	<p>Corrective service diskettes (CSDs) for WebSphere MQ:</p> <p>SPARC systems: 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

¹ This patch is not required for silent installations. For more information about silent installations, see [Appendix B](#).

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, see the following Web site for download and installation information:

<http://www.ibm.com/software/integration/mqfamily/support/summary/sun.html>

3. To continue completing pre-installation tasks, go to the "[Create Required UNIX Groups and User](#)" section on page 2-26.

Checking the Software Requirements on Tru64 UNIX

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: HP Tru64 UNIX V5.1B <hr/> Software Development Kit (SDK) v 1.4.2 for the Tru64 UNIX Operating System for the Java Platform (JDK 1.4.2) <hr/> Operating system subsets: OSFCMPLRS OSFLIBA OSFPGMR OSFSER OSFX11
Oracle DCE Integration (part of Oracle Advanced Security) or Oracle Net protocol support for DCE	DCE v4.3.1 or later: DCECDS DECRTS DCESEC
Oracle Spatial	X Window and X/Motif software: OSFXDEVX OSFXINCLUDE OSFXLIBA Note: This software is required only to build the sample programs.
Oracle Messaging Gateway	IBM MQSeries V5.1 client and server: MQS_CLIENT MQS_SERVER MQSeries classes for Java and MQSeries classes for Java Message Service V5.2.2 (SupportPac MA88): MQS_MQJAVA
PL/SQL native compilation, Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Compaq C Compiler V6.5-207 (dtk): DTCCMLR Compaq C++ Version 6.5-014: CXXBASE CXXLIB CXXOLD
Oracle JDBC/OCI Drivers	You can use the following optional JDK versions with the Oracle JDBC/OCI drivers, however they are not required for the installation: <ul style="list-style-type: none"> ■ SDK v 1.3.1-5 ■ SDK v 1.2.2-12 Note: SDK v 1.4.2 must be installed before you install the software.

Installation Type or Product	Requirement
Oracle Transparent Gateway for Sybase	Either of the following: <ul style="list-style-type: none"> ■ Sybase Adaptive Server Enterprise V12 or 12.5 ■ Sybase Open Client V12 or 12.5 Open Client is required if the Sybase server is not on the local node.

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

1. To determine which version of Tru64 UNIX is installed, enter the following command:

```
# /usr/sbin/sizer -v
Compaq Tru64 UNIX V5.1B (Rev. 2650); Mon Nov 3 10:13:28 PST 200
```

In this example, the version shown is V5.1B. If necessary, see your operating system documentation for information about upgrading the operating system.

2. To determine whether Java SDK 1.4.2 is installed, enter the following command:

```
# /usr/sbin/setld -i JAVA142 | more
```

If Java SDK 1.4.2 is installed, this command displays the paths to all of the installed files. Note the path of the Java home directory. You must specify this value during the installation. The default path is:

```
/usr/opt/java142
```

If this command returns the message Unknown subset, Java SDK 1.4.2 is not installed. Download 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.html>

Check for Required Patches

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

Installation Type or Product	Requirement
All installations	<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, see the following Web site for download and installation information:

<http://www.ibm.com/software/integration/mqfamily/support/summary/dig.html>

Create Required UNIX Groups and User

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

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

- The OSDBA group (`dba`)

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

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

- The OSOPER group (`oper`)

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

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

- An unprivileged user (`nobody` on most platforms)

On platforms other than HP-UX, you must verify that the unprivileged user `nobody` exists on the system. On HP-UX systems, you must create a new unprivileged user. The `nobody` user, or the HP-UX user that you create, must own the external jobs (`ext job`) executable after the installation.

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

- The Oracle Inventory group (`oinstall`)

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

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

- The Oracle software owner user (`oracle`)

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

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

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

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

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

Creating the Oracle Inventory Group

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

Determine Whether the Oracle Inventory Group Exists

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

- AIX, Linux x86, or Linux Itanium:


```
# more /etc/oraInst.loc
```
- Other operating 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, create the Oracle Inventory group as follows, depending on your operating system:

- AIX:
 1. Enter the following command:

```
# smit security
```
 2. Choose the appropriate menu items to create the `oinstall` group.
 3. Press F10 to exit.
- Other operating systems:

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, create it as follows, depending on your operating system. In the following procedure, use the group name `dba` unless a group with that name already exists.

- AIX:
 1. Enter the following command:

```
# smit security
```
 2. Choose the appropriate menu items to create the `dba` group.
 3. Press F10 to exit.
- Other operating systems:

Enter the following command to create the `dba` group:

```
# /usr/sbin/groupadd dba
```

Creating an OSOPER Group (Optional)

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

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

If you require an new OSOPER group, create it as follows, depending on your operating system. In the following procedure, use the group name `oper` unless a group with that name already exists.

- AIX:
 1. Enter the following command:


```
# smit security
```
 2. Choose the appropriate menu items to create the `oper` group.
 3. Press F10 to exit.
- Other operating systems:

Enter the following command to create the `oper` group:

```
# /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 one of the following commands:

- Solaris:


```
# id -a oracle
```
- Other operating systems:


```
# id oracle
```

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

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

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

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

- If you want to use the existing Oracle software owner user, and the user's primary group is the Oracle Inventory group, see one of the following sections:
 - [HP-UX Only: Creating an Unprivileged User](#) on page 2-31
 - [Verifying that the UNIX User nobody Exists](#) on page 2-31

- To modify an existing user, see the "[Modify an Existing Oracle Software Owner User](#)" section on page 2-31.
- 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, depending on your operating system. In the following procedure, use the user name `oracle` unless a user with that name already exists.

- AIX:
 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
```
- Other operating systems:
 1. To create the `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -g oinstall -G dba[,oper] oracle
```

In this command:

- The `-g` option specifies the primary group, which must be the Oracle Inventory group, for example `oinstall`
 - The `-G` option specifies the secondary groups, which must include the OSDBA group and if required, the OSOPER group, for example `dba` or `dba, oper`
2. Set the password of the `oracle` user:

```
# passwd oracle
```

See one of the following sections to continue:

- [HP-UX Only: Creating an Unprivileged User](#) on page 2-31
- [Verifying that the UNIX User `nobody` Exists](#) on page 2-31

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 depending on your operating system:

- AIX:
 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.
- Other operating systems:

Enter a command similar to the following, specifying the primary group using the `-g` option and any required secondary groups using the `-G` option:

```
# /usr/sbin/usermod -g oinstall -G dba[,oper] oracle
```

HP-UX Only: Creating an Unprivileged User

Note: This section applies only to HP-UX systems. If you are installing the software on a different system, see the ["Verifying that the UNIX User nobody Exists"](#) section on page 2-31.

You must create a new unprivileged local user to own the `ext job` executable. You can use any user name for this user, but this document uses the name `ext job`. To create the `ext job` user, enter the following command:

```
# /usr/sbin/useradd extjob
```

To continue, see the ["HP-UX Only: Grant Privileges to the OSDBA Group"](#) section on page 2-32.

Verifying that the UNIX User nobody Exists

Note: This section does not apply to installations on HP-UX 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, enter one of the following commands to create it:

- AIX:

```
# smit security
```

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

- Other operating systems:

```
# /usr/sbin/useradd nobody
```

HP-UX Only: Grant Privileges to the OSDBA Group

Note: If you are installing Oracle Database on a platform other than HP-UX, go to the ["Configure Kernel Parameters and Shell Limits"](#) section on page 2-33.

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 reboots:

```
dba MLOCK RTPRIO RTSCHED
```

3. Save the file and exit from the text editor.
4. Enter the following command to grant the privileges to the OSDBA group:

```
# /usr/sbin/setprivgrp -f /etc/privgroup
```

5. Enter the following command to verify that the privileges are set correctly:

```
# /usr/bin/getprivgrp dba
```

Configure Kernel Parameters and Shell Limits

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

Depending on your operating system, see one of the following sections for information about configuring the system kernel parameters and shell limits:

- [Configuring Shell Limits and System Configuration Parameters on AIX](#) on page 2-33
- [Configuring Kernel Parameters on HP-UX](#) on page 2-34
- [Configuring Kernel Parameters on Linux](#) on page 2-36
- [Configuring Kernel Parameters on Solaris](#) on page 2-39
- [Configuring Kernel Subsystem Attributes on Tru64 UNIX](#) on page 2-40

Configuring Shell Limits and System Configuration Parameters on AIX

Oracle recommends that you set shell limits and system configuration parameters as described in this section.

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.
4. To continue completing pre-installation tasks, go to the ["Identify Required Software Directories"](#) section on page 2-41.

Configuring Kernel Parameters on HP-UX

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. See your operating system documentation for more information about tuning kernel parameters.

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)

Parameter	Recommended Formula or Value
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.
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, see 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 reboot 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, see 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 reboot the system.

- 2. If necessary, when the system restarts, log in and switch user to root.
- 3. To continue completing pre-installation tasks, go to the ["Identify Required Software Directories"](#) section on page 2-41.

Configuring Kernel Parameters on Linux

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
semms	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

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
semnsl, semmns, semopm, and semmni	# /sbin/sysctl -a grep sem This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and shmmni	# /sbin/sysctl -a grep shm
file-max	# /sbin/sysctl -a grep file-max
ip_local_port_range	# /sbin/sysctl -a grep ip_local_port_range This command displays a range of port numbers.

2. If the value of any kernel parameter is different to the recommended value, complete the following steps:
 - a. Using any text editor, create or edit the `/etc/sysctl.conf` file and add or edit lines similar to the following:

Note: Include lines only for the kernel parameter values that you want to change. For the semaphore parameters (`kernel.sem`), you must specify all four values. However, if any of the current values are larger than the recommended value, specify the larger value.

```
kernel.shmall = 2097152
kernel.shmmax = 2147483648
kernel.shmmni = 4096
kernel.sem = 250 32000 100 128
fs.file-max = 65536
net.ipv4.ip_local_port_range = 1024 65000
```

By specifying the values in the `/etc/sysctl.conf` file, they persist when you reboot the system.

- b. Enter the following command to change the current values of the kernel parameters:

```
# /sbin/sysctl -p
```

Review the output from this command to verify that the values are correct. If the values are incorrect, edit the `/etc/sysctl.conf` file, then enter this command again.

- c. On SUSE systems only, enter the following command to cause the system to read the `/etc/sysctl.conf` file when it reboots:

```
# /sbin/chkconfig boot.sysctl on
```

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 <code>limits.conf</code>	Hard Limit
Maximum number of open file descriptors	<code>nofile</code>	65536
Maximum number of processes available to a single user	<code>nproc</code>	16384

To increase the shell limits:

1. Add the following lines to `/etc/security/limits.conf` file:

```
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 (`cs` 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. To continue completing pre-installation tasks, go to the ["Identify Required Software Directories"](#) section on page 2-41.

Configuring Kernel Parameters on Solaris

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_semvmx	32767
shmsys:shminfo_shmmax	4294967295
shmsys:shminfo_shmmin ¹	1
shmsys:shminfo_shmmni	100
shmsys:shminfo_shmseg 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_semmsl=256
set semsys:seminfo_semvmx=32767
set shmsys:shminfo_shmmax=4294967295
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=100
set shmsys:shminfo_shmseg=10
```

- c. Enter the following command to reboot the system:

```
# /usr/sbin/reboot
```

- d. When the system restarts, log in and switch user to `root`.
3. To continue completing pre-installation tasks, go to the ["Identify Required Software Directories"](#) section on page 2-41.

Configuring Kernel Subsystem Attributes on Tru64 UNIX

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.
rdg	msg_size	32768
	max_objs	5120
	max_async_req	256
	max_sessions	500 (or at least 20 plus the value of the <code>PROCESSES</code> 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
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 reboot the system:

```
# /sbin/shutdown -r now
```

- e. When the system restarts, log in and switch user to `root`.

Identify Required Software Directories

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

- Oracle base directory
- Oracle Inventory 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.

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.

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 to view the contents of the `oraInst.loc` file:

- AIX, Linux x86, or Linux Itanium:

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

- Other operating 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`). 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 to view the contents of the `oratab` file:

- Solaris:

```
# more /var/opt/oracle/oratab
```
- Other operating systems:

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

- HP-UX:

```
# bdf oracle_base_path
```
- Other operating systems:

```
# df -k oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, see the "[Create Directories for Oracle Database or Recovery Files](#)" section on page 2-46.

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

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

Creating a New Oracle Base Directory

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

Requirement	Free Disk Space
The Oracle base directory will contain only software files.	Up to 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 -k` command (or `bdf` command on HP-UX) to determine the free disk space on each mounted file system.
2. From the display, identify a file system that has appropriate free space.
The file system that you identify can be a local file system, a cluster file system, or an NFS file system on a certified NAS device.
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. When you are configuring the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify this directory.

Choose a Storage Option for Oracle Database and Recovery Files

The following table shows the storage options supported for storing 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.

For all installations, you must choose the storage option that you want to use for 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.

Storage Option	File Types Supported	
	Database	Recovery
File system	Yes	Yes
Automatic Storage Management	Yes	Yes
Raw devices	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.
- Oracle recommends that you choose ASM as the storage option for database and recovery files.
- For more information about these storage options, see the "[Database Storage Options](#)" section on page 1-5.

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

- To use a file system for database or recovery file storage, see the "[Create Directories for Oracle Database or Recovery Files](#)" section on page 2-46.
- To use ASM for database or recovery file storage, see the "[Configure Disks for Automatic Storage Management](#)" section on page 2-49.
- To use raw devices (partitions or logical volumes) for database file storage, see the "[Configure Raw Devices](#)" section on page 2-66.

Create Directories for Oracle Database or Recovery Files

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

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:

- The default path suggested by the Installer for the database file directory is a subdirectory of the Oracle base directory.

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 allows you to distribute physical I/O and create separate control files on different devices for increased reliability. It also allows you to fully implement the OFA guidelines described in Appendix D, "Optimal Flexible Architecture". 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.

- The file system that you choose should have at least 2 GB of free disk space.

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

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

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

- The default path suggested by the Installer for the flash recovery area is a subdirectory of the Oracle base directory.
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 database or recovery files on a separate file system to the Oracle base directory.

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

1. If necessary, configure the file systems that you want to use and mount them on each node.
2. Use the `df -k` command (or `bdf` command on HP-UX) 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
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:
 - Database file directory:


```
# mkdir /mount_point/oradata
# chown oracle:oinstall /mount_point/oradata
# chmod 775 /mount_point/oradata
```
 - Recovery file directory (flash recovery area):


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

6. If you also want to use ASM or raw devices for storage, see one of the following sections:
 - [Configure Disks for Automatic Storage Management](#) on page 2-49
 - [Configure Raw Devices](#) on page 2-66Otherwise, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

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. See [Appendix C](#) 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.

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

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 2-53.

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:

- Solaris:

```
# more /var/opt/oracle/oratab
```

- Other operating systems:

```
# 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 AIX](#) on page 2-53
- [Configuring Disks for ASM on HP-UX](#) on page 2-54
- [Configuring Disks for ASM on Linux](#) on page 2-55
- [Configuring Disks for ASM on Solaris](#) on page 2-63
- [Configuring Disks for ASM on Tru64 UNIX](#) on page 2-65

Configuring Disks for ASM on AIX

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

1. If necessary, install the disks that you intend to use for the ASM disk group and reboot the system.
2. To make sure that the disks are available, enter the following command:

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

The output from this command is similar to the following:

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

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

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

4. Enter the following command 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.

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

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

6. Enter commands similar to the following to change the owner, group, and permissions on the character 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.

7. If you also want to use raw devices for storage, see the ["Configure Raw Devices"](#) section on page 2-66.

Otherwise, see the ["Stop Existing Oracle Processes"](#) section on page 2-92.

Configuring Disks for ASM on HP-UX

To configure disks for use with ASM on HP-UX, follow these steps:

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

```
# /usr/sbin/ioscan -fun -C disk
```

The output from this command is similar to the following:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	0	0/0/1/0.6.0	sdisk	CLAIMED	DEVICE	HP DVD-ROM 6x/32x
			/dev/dsk/c0t6d0		/dev/rdisk/c0t6d0	
disk	1	0/0/1/1.2.0	sdisk	CLAIMED	DEVICE	SEAGATE ST39103LC
			/dev/dsk/c1t2d0		/dev/rdisk/c1t2d0	

This command displays information about each disk attached to the system, including the block device name (/dev/dsk/cxydz) and the character raw device name (/dev/rdisk/cxydz).

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 a disk group, enter the following command to verify that it is not already part of an LVM volume group:

```
# /sbin/pvdisplay /dev/dsk/cxydz
```

If this command displays volume group information, the disk is already part of a volume group. The disks that you choose must not be part of an LVM volume group.

Note: If you are using different volume management software, for example VERITAS Volume Manager, see the appropriate documentation for information about verifying that a disk is not in use.

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

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

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

6. If you also want to use raw devices for storage, see the "[Configure Raw Devices](#)" section on page 2-66.

Otherwise, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

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 Database CD-ROM in the `RPMS/asmlib` directory and on the DVD-ROM in the `db/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
oracleasm-lib-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 \
    oracleasm-lib-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 \
    oracleasm-lib-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.

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 disk devices that you intend to use for the disk group and reboot 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	/dev/hdxn	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, /dev/hda is the first disk on the first IDE bus.
SCSI disk	/dev/sdxn	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, /dev/sda is the first disk on the first SCSI bus.
RAID disk	/dev/rd/cxdypz /dev/ida/cxdypz	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <i>x</i> is a number that identifies the controller, <i>y</i> is a number that identifies the disk, and <i>z</i> is a number that identifies the partition. For example, /dev/ida/c0d1 is the second logical drive on the first controller.

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

Note: 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 DASDs that you intend to use for the disk group and reboot 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 ld1` 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.

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 rebooting 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 disk devices that you intend to use for the disk group and reboot 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	/dev/hdxn	In this example, <i>x</i> is a letter that identifies the IDE disk and <i>n</i> is the partition number. For example, /dev/hda is the first disk on the first IDE bus.
SCSI disk	/dev/sdxn	In this example, <i>x</i> is a letter that identifies the SCSI disk and <i>n</i> is the partition number. For example, /dev/sda is the first disk on the first SCSI bus.
RAID disk	/dev/rd/cxdypz /dev/ida/cxdypz	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <i>x</i> is a number that identifies the controller, <i>y</i> is a number that identifies the disk, and <i>z</i> is a number that identifies the partition. For example, /dev/ida/c0d1 is the second logical drive on the first controller.

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

Note: 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 DASDs that you intend to use for the disk group and reboot 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 to bind the disk devices to raw devices:

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

```
# /usr/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.

- 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 reboots.

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

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

```
# /usr/sbin/raw -qa
```

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

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

- 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 reboots, enter the following command:

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

5. If you also want to use raw devices for storage, see the "[Configure Raw Devices](#)" section on page 2-66.

Otherwise, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

Configuring Disks for ASM on Solaris

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

1. If necessary, install the disks that you intend to use for the disk group and reboot 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, see the appropriate documentation for information about determining which disk devices it is managing.

```
# vxdiskconfig
# /usr/sbin/vxdisk list
```

If this command displays disk group information associated with a disk device, the disk is already part of 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. If you plan to use existing slices, enter the following command 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.

5. Enter commands similar to the following 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.

6. If you also want to use raw devices for storage, see the ["Configure Raw Devices"](#) section on page 2-66.

Otherwise, see the ["Stop Existing Oracle Processes"](#) section on page 2-92.

Configuring Disks for ASM on Tru64 UNIX

To configure disks for use with ASM on Tru64 UNIX, follow these steps:

1. If necessary, install the disks that you intend to use for the disk group and reboot the system.
2. To make sure that the disks are available, enter the following command:

```
# file /dev/rdisk/dsk*c
```

The output from this command is similar to the following:

```
/dev/rdisk/dsk0c:      character special (19/38) SCSI #1 "RZ1CB-CA" disk #2
                    (SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk10c:   character special (19/218) SCSI #2 "RZ1DB-CS" disk #4
                    (SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk11c:   character special (19/234) SCSI #2 "RZ1DB-CS" disk #5
                    (SCSI ID #3) (SCSI LUN #0)
/dev/rdisk/dsk12c:   character special (19/250) SCSI #2 "RZ1DB-CS" disk #6
                    (SCSI ID #4) (SCSI LUN #0)
```

This command displays information about each disk attached to the system, including the raw character device name for partition *c*, which represents the entire disk (*/dev/rdisk/dsknc*).

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/rdisknc  
# chmod 660 /dev/rdisk/rdisknc
```

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.

9. To continue, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

Configure Raw Devices

Depending on your operating system, see one of the following sections for information about configuring raw devices:

- [Configuring Raw Logical Volumes on AIX](#) on page 2-67
- [Configuring Raw Logical Volumes on HP-UX](#) on page 2-71
- [Configuring Raw Partitions or Raw Logical Volumes on Linux](#) on page 2-75
- [Configuring Raw Partitions or Raw Logical Volumes on Solaris](#) on page 2-84

Configuring Raw Logical Volumes on AIX

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

Create a Volume Group

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

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

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

The output from this command is similar to the following:

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

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

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

4. Enter the following command 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:

```
# 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.

- 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 \  
-C PhysicalVolumes
```

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

Command Option	SMIT Field	Sample Value and Description
-y <i>VGname</i>	VOLUME GROUP name	oracle_vg1 Specify the name for the volume group. The name that you specify could be a generic name, as shown, or 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 allows 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.
-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.

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 2-5](#) lists the number and size of the logical volumes that you must create for database files.

Table 2–5 Raw Logical Volumes Required for Database Files on AIX

Number	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	500	UNDOTBS1 tablespace: <i>dbname_undotbs1_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	120	Two online redo log files (where <i>m</i> is the log number, 1 or 2): <i>dbname_redo1_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>

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:

```
# /usr/sbin/mklv -y test_sysaux_raw_500m -T 0 -w n -s n -r n oracle_vg1 16
```

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/rdbname*
# chmod 660 /dev/rdbname*
```

Create the DBCA Raw Device Mapping File

To allow 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:

```
system=/dev/rdbname_system_raw_500m
sysaux=/dev/rdbname_sysaux_raw_500m
example=/dev/rdbname_example_raw_160m
users=/dev/rdbname_users_raw_120m
temp=/dev/rdbname_temp_raw_250m
undotbs1=/dev/rdbname_undotbs1_raw_500m
redo1_1=/dev/rdbname_redo1_1_raw_120m
redo1_2=/dev/rdbname_redo1_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/rrac_redo1_1_raw_120m
```

- For a single-instance database, the file must specify one automatic undo tablespace datafile (`undotbs1`), and at least two redo log files (`redo1_1`, `redo1_2`).
 - 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 datafile.
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.
 8. To continue, see the "Stop Existing Oracle Processes" section on page 2-92.

Configuring Raw Logical Volumes on HP-UX

The following subsections describe how to configure raw logical volumes on HP-UX.

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. See the HP-UX documentation for more information about using SAM.

Create a Volume Group

To create a volume group, follow these steps:

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

```
# /sbin/ioscan -fun -C disk
```

The output from this command is similar to the following:

```
Class I H/W Path Driver S/W State H/W Type Description
=====
disk 0 0/0/1/0.6.0 sdisk CLAIMED DEVICE HP DVD-ROM 6x/32x
/dev/dsk/c0t6d0 /dev/rdisk/c0t6d0
disk 1 0/0/1/1.2.0 sdisk CLAIMED DEVICE SEAGATE ST39103LC
/dev/dsk/c1t2d0 /dev/rdisk/c1t2d0
disk 2 0/0/2/0.2.0 sdisk CLAIMED DEVICE SEAGATE ST118202LC
/dev/dsk/c2t2d0 /dev/rdisk/c2t2d0
```

This command displays information about each disk attached to the system, including the block device name (`/dev/dsk/cxydz`) and the character raw device name (`/dev/rdisk/cxydz`).

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 to verify that it is not already part of an LVM volume group:

```
# /sbin/pvdisplay /dev/dsk/cxytdz
```

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:

```
# /sbin/pvcreate /dev/rsk/cxytdz
```

6. 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
```

7. To identify used device minor numbers, enter the following command:

```
# ls -la /dev/*/group
```

This command displays information about the device numbers used by all configured volume groups, similar to the following:

```
crw-r----- 1 root  sys      64 0x000000 Mar  4  2002 /dev/vg00/group
crw-r--r--  1 root  sys      64 0x010000 Mar  4  2002 /dev/vg01/group
```

In this example, 64 is the major number used by all volume group devices and 0x000000 and 0x010000 are the minor numbers used by volume groups `vg00` and `vg01` respectively. Minor numbers have the format `0xnn0000`, where `nn` is a number in the range 00 to the value of the `maxvgs` kernel parameter minus 1. The default value for the `maxvgs` parameter is 10, so the default range is 00 to 09.

8. Identify an appropriate minor number that is unused.
9. 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
- `0xnn0000` is a minor number that is unused
- `/dev/dsk/cxytdz . . .` is a list of one or more block device names for the disks that you want to add to the volume group

Create Raw Logical Volumes in the New Volume Group

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

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

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

2. Identify the logical volumes that you must create.

[Table 2-6](#) lists the number and size of the logical volumes that you must create for database files.

Table 2-6 Raw Logical Volumes Required for Database Files on HP-UX

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	500	SYSTEM tablespace: <code>dbname_system_raw_500m</code>
1	500	SYSAUX tablespace: <code>dbname_sysaux_raw_500m</code>
1	500	UNDOTBS1 tablespace: <code>dbname_undotbs1_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	120	Two online redo log files (where <i>m</i> is the log number, 1 or 2): <code>dbname_redo1_m_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>

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 a 500 MB logical volume in the `oracle_vg` volume group for the SYSAUX tablespace of a database named `test`:

```
# /sbin/lvcreate -n test_sysaux_raw_500m -L 500 /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*
```

Create the DBCA Raw Device Mapping File

To allow 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:

```
system=/dev/vg_name/rdbname_system_raw_500m
sysaux=/dev/vg_name/rdbname_sysaux_raw_500m
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
redo1_1=/dev/vg_name/rdbname_redo1_1_raw_120m
redo1_2=/dev/vg_name/rdbname_redo1_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/rrac_redo1_1_raw_120m
```

- For a single-instance database, the file must specify one automatic undo tablespace datafile (*undotbs1*), and at least two redo log files (*redo1_1*, *redo1_2*).
 - 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 datafile.
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.
 8. To continue the installation, see the ["Stop Existing Oracle Processes"](#) section on page 2-92.

Configuring Raw Partitions or Raw Logical Volumes on Linux

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

Choose Between Raw Logical Volumes and Raw Partitions

The procedures in the following sections describe how to create raw partitions or raw logical volumes for Oracle database file storage. You can use raw partitions for database file storage on all supported Linux distributions.

Alternatively, on Red Hat Enterprise Linux 3 and SUSE Linux Enterprise Server you can use Logical Volume Manager (LVM) to create logical volumes of the required size and use them instead of raw partitions. LVM can provide benefits when used for storing Oracle database files on raw devices. For example, LVM enables you to overcome the limitations on the number of partitions that you can create on a single physical disk.

Identify the Required Raw Partitions

Table 2-7 lists the number and size of the raw partitions or logical volumes that you must configure for database 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 2-7 Raw Partitions or Logical Volumes Required for Database Files on Linux

Number	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	500	UNDOTBS1 tablespace: <i>dbname_undotbs1_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	120	Two online redo log files (where <i>m</i> is the log number, 1 or 2): <i>dbname_redo11_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 UNDOTBS1 raw device.

To create the required raw partitions or logical volumes, see one of the follow sections:

- [Creating the Required Raw Partitions on IDE, SCSI, or RAID Devices](#) on page 2-76
- [Creating the Required Raw Logical Volumes](#) on page 2-78
- [zSeries Linux Only: Creating the Required Raw Partitions on DASDs](#) on page 2-80

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 disk devices that you intend to use for the raw partitions and reboot the system.

Note: Because the number of partitions that you can create on a single device is limited, you might need to create the required raw partitions on more than one device.

2. To identify the device name for the disks that you want to use, 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, see the ["Bind the Partitions or Logical Volumes to Raw Devices"](#) section on page 2-81.

Creating the Required Raw Logical Volumes

For single instance databases, you can use raw logical volume manager (LVM) volumes for Oracle database file storage. To configure the required raw logical volumes, follow these steps:

1. On zSeries Linux only, if you intend to use direct access storage devices (DASDs) for the volume group, follow these steps:
 - a. If necessary, install or configure the DASDs that you intend to use for the volume group and reboot 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.

Alternatively, you could use the `-d 1d1` 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 IDE or SCSI devices in the volume group, follow these steps:
- a. If necessary, install or configure the disk devices that you intend to use for the volume group and reboot 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.

- 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. To continue, see the ["Bind the Partitions or Logical Volumes to Raw Devices"](#) section on page 2-81.

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 DASDs that you intend to use for the raw partitions and reboot 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 2-81.

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.

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. 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:

- Red Hat:
 1. To determine what raw devices are already bound to other devices, enter the following command:

```
# /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.

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 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
```

4. 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 reboots.

■ SUSE:

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

```
# /usr/sbin/raw -qa
```

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

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

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 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
```

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

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

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

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

Create the DBCA Raw Device Mapping File

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

- 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
```

- 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.

- Change directory to the `$ORACLE_BASE/oradata/dbname` directory.
- Edit the `dbname_raw.conf` file in any text editor to create a file similar to the following:

```
system=/dev/raw/raw1
sysaux=/dev/raw/raw2
example=/dev/raw/raw3
users=/dev/raw/raw4
temp=/dev/raw/raw5
undotbs1=/dev/raw/raw6
redo1_1=/dev/raw/raw8
redo1_2=/dev/raw/raw9
control1=/dev/raw/raw12
control2=/dev/raw/raw13
spfile=/dev/raw/raw14
pwdfile=/dev/raw/raw15
```

Use the following guidelines when creating or editing this file:

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

```
database_object_identifier=raw_device_path
```

- For a single-instance database, the file must specify one automatic undo tablespace datafile (`undotbs1`), and at least two redo log files (`redo1_1`, `redo1_2`).
- 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 datafile.

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.
7. To continue, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

Configuring Raw Partitions or Raw Logical Volumes on Solaris

The following subsections describe how to configure raw partitions or raw logical volumes on Solaris:

- [Configuring Raw Partitions for Oracle Database File Storage](#) on page 2-84
Follow the procedures described in this section if you want to use raw partitions for Oracle database file storage. You can use either raw disk slices or Solaris Volume Manager soft partitions to store the required database files.
-
- Note:** Solaris Volume Manager volumes or soft partitions are supported only for single-instance database file storage. Solaris Volume Manager is not currently supported as a storage option for RAC installations.
-
- [SPARC Only: Configuring Raw Logical Volumes for Oracle Database File Storage](#) on page 2-88
Follow the procedures described in this section only if you are using VERITAS Volume Manager (VxVM) on SPARC systems.

Configuring Raw Partitions for Oracle Database File Storage

This section describes how to configure raw partitions for Oracle database files.

A partition is a slice defined on a disk or on a disk array volume. It can also be a soft partition created using Solaris Volume Manager.

[Table 2-8](#) lists the number and size of the raw partitions that you must configure for database 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 2-8 *Raw Partitions Required for Database Files on Solaris*

Number	Size (MB)	Purpose and Database Object Identifier
1	500	SYSTEM tablespace: system
1	500	SYSAUX tablespace: sysaux
1	500	UNDOTBS1 tablespace: undotbs1

Table 2–8 (Cont.) Raw Partitions Required for Database Files on Solaris

Number	Size (MB)	Purpose and Database Object Identifier
1	250	TEMP tablespace: temp
1	160	EXAMPLE tablespace: example
1	120	USERS tablespace: users
2	120	Two online redo log files (where <i>m</i> is the log number, 1 or 2): redo1_ <i>m</i>
2	110	First and second control files: control{1 2}
1	5	Server parameter file (SPFILE): spfile
1	5	Password file: pwdfile

To configure raw partitions for database files, 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. If necessary, install or configure the disks that you intend to use and reboot the system.
3. If you want to use Solaris Volume Manager soft partitions, see the Solaris Volume Manager documentation for information about how to create them.
The previous table shows the number and size of the partitions that you require.
4. 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 database files.

Note: To prevent the 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.
5. If you plan to use existing partitions, enter the following command 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. The device name for a Solaris Volume Manager partition is similar to `/dev/md/dsk/dnnn`, where `dnnn` is the soft partition name.

6. Enter commands similar to the following to change the owner, group, and permissions on the character raw device file for each partition:

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.

- Solaris Volume Manager soft partitions:

```
# chown oracle:dba /dev/md/rdisk/d100
# chmod 660 /dev/md/rdisk/d100
```

- Disk slices:

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


7. Follow these steps to create 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 that specifies Solaris Volume Manager soft partitions. If you are using disk slices, specify the appropriate device file name for each slice, for example:

```
/dev/rdisk/cxdyztzsn
```

```
system=/dev/md/rdisk/d100
sysaux=/dev/md/rdisk/d101
example=/dev/md/rdisk/d102
users=/dev/md/rdisk/d103
temp=/dev/md/rdisk/d104
undotbs1=/dev/md/rdisk/d105
redo1_1=/dev/md/rdisk/d106
redo1_2=/dev/md/rdisk/d107
control1=/dev/md/rdisk/d108
control2=/dev/md/rdisk/d109
spfile=/dev/md/rdisk/d110
pwdfile=/dev/md/rdisk/d111
```

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 previous table shows the database object identifier for each database file.

- Specify one automatic undo tablespace datafile (`undotbs1`) and two redo log files (`redo1_1`, `redo1_2`).
 - 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.
8. To continue, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

SPARC Only: Configuring Raw Logical Volumes for Oracle Database File Storage

This section describes how to configure raw logical volumes using VERITAS Volume Manager (VxVM) on SPARC systems.

Create a Disk Group

To create a disk group, follow these steps:

1. If necessary, install the disks that you intend to use for the disk group and reboot the system.
2. To make sure that the disks are available, enter the following command:

```
# /usr/sbin/format
```

The output from this command is similar to the following:

```
AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <ST34321A cyl 8892 alt 2 hd 15 sec 63>
     /pci@1f,0/pci@1,1/ide@3/dad@0,0
  1. c1t5d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
     /pci@1f,0/pci@1/scsi@1/sd@5,0
```

This command displays information about each disk attached to the system, including the device name (`cxydz`).

3. From the list, identify the device names for the disk devices that you want to add to a disk group, then press `Ctrl/D` to exit from the `format` utility.
4. Enter the following command to verify that the devices you identified are not mounted as file systems:

```
# df -k
```

This command displays information about the partitions (slices) on disk devices that are mounted as file systems. The device name for a slice includes the disk device name followed by the slice number, for example `cxydzsn`, where `sn` is the slice number. Slice 2 (`s2`) represents the entire disk. The disk devices that you choose must not be shown as mounted partitions.

- 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, see the appropriate documentation for information about determining which disk devices it is managing.

```
# /usr/sbin/vxdiskconfig
# /usr/sbin/vxdisk list
```

The `vxdisk list` command identifies the disk devices that are already configured in a disk group. The word `online` in the `STATUS` column also identifies disks that have been initialized and placed under VxVM control. The word `error` in the `STATUS` column identifies disks that are not initialized.

The disk devices that you choose must not be in an existing disk group.

- If the disk devices that you want to use are not initialized, enter a command similar to the following to initialize each disk:

```
# /usr/sbin/vxdiskadd cxytdz
```

- To create a 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 init diskgroup diskname=deviceName ...
```

In this example:

- `diskgroup` is the name of the disk group that you want to create, for example, `oradg`
- `diskname` is an administrative name that you assign to a disk, for example `orad01`
- `deviceName` is the device name, for example, `c1t0d0`

Create Raw Logical Volumes in the New Disk Group

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

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

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

- Identify the logical volumes that you must create.

[Table 2–9](#) lists the number and size of the logical volumes that you must create for database files.

Table 2–9 Raw Logical Volumes Required for Database Files on Solaris

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	500	SYSTEM tablespace: <code>dbname_system_raw_500m</code>
1	500	SYSAUX tablespace: <code>dbname_sysaux_raw_500m</code>

Table 2–9 (Cont.) Raw Logical Volumes Required for Database Files on Solaris

Number	Size (MB)	Purpose and Sample Logical Volume Name
1	500	UNDOTBS1 tablespace: <i>dbname_undotbs1_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	120	Two online redo log files (where <i>m</i> is the log number, 1 or 2): <i>dbname_redo1_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>

3. To create the 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 a 500 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_500m 500m \  
user=oracle group=dba mode=660
```

Create the DBCA Raw Device Mapping File

To allow 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:

```
system=/dev/vx/rdsk/diskgroup/dbname_system_raw_500m
sysaux=/dev/vx/rdsk/diskgroup/dbname_sysaux_raw_500m
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
redo1_1=/dev/vx/rdsk/diskgroup/dbname_redo1_1_raw_120m
redo1_2=/dev/vx/rdsk/diskgroup/dbname_redo1_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 single-instance database, the file must specify one automatic undo tablespace datafile (`undotbs1`), and at least two redo log files (`redo1_1`, `redo1_2`).
 - 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 datafile.
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.
 8. To continue, see the "[Stop Existing Oracle Processes](#)" section on page 2-92.

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, see the "[Configure the oracle User's Environment](#)" section on page 2-93 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.

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`, `ORACLE_BASE`, and `ORACLE_SID` 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. On AIX only, add one of the following lines to the file to set the AIXTHREAD_SCOPE environment variable to S (system-wide thread scope):

- Bourne, Bash, or Korn shell:

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

- C shell:

```
setenv AIXTHREAD_SCOPE S
```

Note: For more information about this environment variable, see 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:

- Bash shell on Red Hat Enterprise Linux:

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

- 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:

- a. Use the `df -k` command (or `bdf` command on HP-UX) 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
```

14. Enter commands similar to the following to set the ORACLE_BASE and ORACLE_SID environment variables:

■ Bourne, Bash, or Korn shell:

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

■ C shell:

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

In these examples, `/u01/app/oracle` is the Oracle base directory that you created or identified earlier and `sales` is the name that you want to call the database (typically no more than five characters).

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.

Installation Tasks

This chapter describes how to start the Oracle Universal Installer and install Oracle products on your system. It includes information about the following topics:

Note: Review the information in [Chapter 1, "Installation Overview"](#) and complete the tasks listed in [Chapter 2, "Pre-installation Tasks"](#) before beginning the installation.

- [Installation Overview](#)
- [Download Oracle Software from the OTN Web Site](#)
- [Copy the Oracle Database Software to a Hard Disk](#)
- [AIX Only: Run the rootpre.sh Script](#)
- [Install the Oracle Database Software](#)

Installation Overview

The Oracle Database software is available on disc (CD-ROM or DVD-ROM) or you can download it from the Oracle Technology Network (OTN) Web site. In most cases, you use the graphical user interface (GUI) provided by the Oracle Universal Installer to install the software. However, you can also use the Installer to complete silent-mode installations, without using the GUI. See [Appendix B](#) for information about silent-mode installations.

This chapter describes how to install Oracle Database from the product media or from the hard disk. To install the software from the hard disk, you must either download it from OTN and unpack it, or copy it from the discs, if you have them. See one of the following sections, depending on the method that you want to use:

Note: If the Oracle Database 10g software for your platform is on two discs, you can avoid having to change discs during the installation by running the Installer from the DVD-ROM or by copying the contents of the installation discs to a hard disk.

- To copy the software to a hard disk, see the ["Copy the Oracle Database Software to a Hard Disk"](#) section on page 3-3.
- To install the software from disc or from an existing hard disk location, see one of the following sections, depending on your platform:
 - AIX:
See the ["AIX Only: Run the rootpre.sh Script"](#) section on page 3-5.
 - Other platforms:
See the ["Install the Oracle Database Software"](#) section on page 3-5.
- To download the software from OTN, see the following section.

Download Oracle Software from the OTN Web Site

This section describes how to download the installation archive files and extract them on your hard disk.

Downloading the Installation Archive Files

To download the installation archive files from OTN:

1. Use any browser to access the software download page on OTN:
`http://www.oracle.com/technology/software/`
2. Navigate to the download page for the product that you want to install.
3. On the download page, identify the required disk space by adding the file sizes for each required file.
The file sizes are listed next to the filenames.
4. Select a file system with enough free space to store and expand the archive files.
In most cases, the available disk space must be at least twice the size of all of the archive files.
5. On the file system that you selected in step 4, create a parent directory for each product, for example `OraDB10g`, to hold the installation directories.

6. Download all of the installation archive files to the directory that you created in step 5.
7. Verify that the files you downloaded are the same size as the corresponding files on OTN.

Extracting the Installation Files

To extract the installation archive files, perform the following steps:

1. If necessary, change directory to the directory that contains the downloaded installation archive files.
2. To uncompress each file, enter a command similar to the following:

```
$ gunzip filename.cpio.gz
```

This command creates files with names similar to the following:

```
filename.cpio
```

3. To extract the installation files, enter commands similar to the following:

```
$ cpio -idmv < filename.cpio
```

Note: See the download page for information about the correct options to use with the `cpio` command on your platform.

Some browsers uncompress files while downloading them, but leave the `.gz` file extension. If these steps do not work, remove the `.gz` extension from the files and repeat step 3.

For each file, this command creates a subdirectory named `Diskn`, where `n` is the disk number identified in the filename.

When you have extracted all of the required installation files, see one of the following sections, depending on your platform:

- AIX:
See the "[AIX Only: Run the rootpre.sh Script](#)" section on page 3-5.
- Other platforms:
See the "[Install the Oracle Database Software](#)" section on page 3-5.

Copy the Oracle Database Software to a Hard Disk

If the software for Oracle Database 10g requires two discs on your platform, you can avoid having to mount and unmount discs during the installation by copying the contents of each CD-ROM to a hard disk on the system, or by installing the software from the DVD-ROM.

Note: If your system does not have a CD-ROM or DVD-ROM drive, you can copy the software from the discs to a file system on another system, then either mount that file system using NFS, or use FTP to copy the files to the system where you want to install the software.

To copy the contents of the discs to a hard disk:

1. Create a single directory on the hard disk to hold the Oracle software:

```
$ mkdir OraDb10g
```

2. Change directory to the directory you created in step 1:

```
$ cd OraDb10g
```

3. Create subdirectories for each installation disc with names similar to the following, where *n* is the disc number:

```
$ mkdir Diskn
```

4. Mount the first disc, if it is not already mounted.

Some platforms automatically mount the disc when you insert it into the drive. If the disc does not mount automatically, see [Appendix A](#) for platform-specific information about mounting it.

5. Copy the contents of the mounted disc to the corresponding new subdirectory as follows:

```
$ cp -R /directory_path Diskn
```

In this example, *n* is the number of the installation CD-ROM and */directory_path* is the CD-ROM mount point directory or the path of the db directory on the DVD-ROM. The following table lists typical mount point directories for each platform:

Platform	Mount Point
AIX, Tru64 UNIX	/cdrom
HP-UX	/SD_CDROM
Linux (Red Hat)	/mnt/cdrom
Linux (SUSE)	/media/cdrom
Solaris (mounted automatically)	/cdrom/cdrom0
Solaris (mounted manually)	/cdrom

6. If necessary, mount the next disc and repeat step 5.

After you have copied all of the required discs, see one of the following sections, depending on your platform:

- AIX:
See the "[AIX Only: Run the rootpre.sh Script](#)" section on page 3-5.
- Other platforms:
See the "[Install the Oracle Database Software](#)" section on page 3-5.

AIX Only: Run the rootpre.sh Script

Note: Do not run the `rootpre.sh` script if you have a later release of the Oracle Database software already installed on this system.

On AIX systems, run the `rootpre.sh` script:

1. Switch user to `root`:

```
$ su - root
```

2. Complete one of the following steps, depending on the location of the installation files:

- If the installation files are on disc, enter a command similar to the following, where *directory_path* is the CD-ROM mount point directory or the path of the `db` directory on the DVD-ROM:

```
# /directory_path/rootpre.sh
```

- If the installation files are on the hard disk, change directory to the `Disk1` directory and enter the following command:

```
# ./rootpre.sh
```

3. Exit from the `root` account:

```
# exit
```

Install the Oracle Database Software

The following sections describe how to install the Oracle software:

- [Reviewing Product-Specific Installation Guidelines](#)
- [Running the Oracle Universal Installer](#)

Reviewing Product-Specific Installation Guidelines

Review the following guidelines before starting the Installer:

- Oracle Universal Installer

Do not use the Oracle Universal Installer from an earlier Oracle product release to install components from this release.

- Products Requiring a Custom Installation

To install the following products, you must choose the Custom installation type:

- Oracle Label Security

Note: Before you perform a custom installation, make sure all databases in the Oracle home where you want to install Oracle Label Security are shut down.

If you want to configure OLS to use Oracle Internet Directory (OID), you must choose the OID option when running the DBCA.

- Oracle Connection Manager
- Oracle Procedural and Transparent Gateways
- Data Mining Scoring Engine
- Installations on a Cluster

If Oracle Cluster Ready Services (CRS) and Oracle Real Application Clusters (RAC) are already installed on the system, the Installer displays the Specify Hardware Cluster Installation Mode screen. You must select **Local Installation** on this screen, unless you want to install Oracle Real Application Clusters.

If you do want to install Oracle Real Application Clusters, see the *Oracle Real Application Clusters Installation and Configuration Guide* instead of using this guide.
- Re-installing Oracle Software

If you re-install Oracle software into an Oracle home directory where Oracle Database is already installed, you must also re-install any components, such as Oracle Partitioning, that were installed before you began the re-installation.

Running the Oracle Universal Installer

Start the Installer and install the software, as follows:

1. If you are installing the software from disc, mount the first disc, if it is not already mounted.

Some platforms automatically mount the disc when you insert it into the drive. If the disc does not mount automatically, see [Appendix A](#) for platform-specific information about mounting it.
2. If necessary, log in as the Oracle software owner user (`oracle`) and set the `DISPLAY` environment variable.
3. To start the Installer, complete one of the following steps depending on the location of the installation files:
 - If the installation files are on disc, enter commands similar to the following, where `directory_path` is the CD-ROM mount point directory or the path of the `db` directory on the DVD-ROM:

```
$ cd /tmp
$ /directory_path/runInstaller
```
 - If the installation files are on the hard disk, change directory to the `Disk1` directory and enter the following command:

```
$ ./runInstaller
```

If the Installer does not appear, see the "[X Windows Display Errors](#)" section on page F-2 for information about troubleshooting.
4. Use the following guidelines to complete the installation:
 - Follow the instructions displayed in the Installer windows. If you need additional information, click **Help**.
 - When the Installer prompts you to run a script with `root` privileges, enter a command similar to the following in a terminal where you are logged in as the `root` user, then click **Continue** or **OK**:

```
# /script_path/script_name
```


- When prompted for the SYS, SYSTEM, SYSMAN, and DBSNMP passwords, Oracle recommends that you specify passwords that:
 - Are at least four characters long
 - Are not the same as the user names
 - Have at least one alphabetic, one numeric, and one punctuation mark character
 - Are not simple or obvious words, such as welcome, account, database, or user

Note: You must remember the passwords that you specify.

- On Tru64 UNIX, when prompted for the JDK home directory, specify the location of JDK 1.4.2 on your system.

To determine the JDK home directory, enter the following command:

```
# /usr/sbin/setld -i JAVA142 | more
```

The default location is `/usr/opt/java142`.

Caution: The Oracle Universal Installer automatically installs the Oracle-supplied version of the Java Runtime Environment (JRE). This version is required to run Oracle Universal Installer and several Oracle assistants. Do not modify the JRE except by using a patch provided by Oracle Support Services.

- On some platforms, if you run the Installer from CD-ROM, you must change discs during the installation. To change discs, unmount the current disc and mount the next disc. See [Appendix A](#) for platform-specific information about unmounting and mounting discs.
 - If you encounter errors while installing or linking the software, see [Appendix F](#) for information about troubleshooting.
5. If you chose an installation type that runs DBCA and Oracle Net Configuration Assistant (NetCA) in interactive mode, you must provide detailed information about configuring your database and network.

If you need assistance when using the DBCA or NetCA in interactive mode, click **Help** on any screen.

Note: If you chose a default installation, the DBCA and NetCA run non-interactively.

6. When all of the configuration tools have finished, click **Exit**, then click **Yes** to exit from the Installer.
7. See [Chapter 4, "Post-installation Tasks"](#) for information about tasks that you must complete after you have installed the software.

Post-installation Tasks

This chapter describes how to complete post-installation tasks after you have installed the software. It includes information about the following topics:

- [Required Post-installation Tasks](#)
- [Recommended Post-installation Tasks](#)
- [Required Product-Specific Post-installation Tasks](#)
- [Installing Oracle Database 10g Products from the Companion CD](#)

You must perform the tasks listed in the "[Required Post-installation Tasks](#)" section. Oracle recommends that you perform the tasks listed in the "[Recommended Post-installation Tasks](#)" section after all installations.

If you installed and intend to use any of the products listed in the "[Required Product-Specific Post-installation Tasks](#)" section, you must perform the tasks listed in the product-specific subsections.

Note: This chapter describes basic configuration only. See the *Oracle Database Administrator's Reference for UNIX Systems* and product-specific administration and tuning guides for more sophisticated configuration and tuning information.

Required Post-installation Tasks

You must perform the tasks described in the following sections after completing an installation:

- [Downloading and Installing Patches](#)
- [Configuring Oracle Products](#)
- [HP-UX Only: Changing Permissions on the extjob Executable](#)

Downloading and Installing Patches

Check 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, click **Register for MetaLink!** and follow the registration instructions.

3. On the main Oracle*Metalink* page, click **Patches**.
4. Select **Simple Search**.
5. Specify the following information, then click **Go**:
 - In the **Search By** field, choose Product or Family, then specify RDBMS Server.
 - In the **Release** field, specify the current release number.
 - In the **Patch Type** field, specify Patchset/Minipack.
 - In the **Platform or Language** field, select your platform.

Configuring Oracle Products

Many Oracle products and options must be configured before you use them for the first time. Before using individual Oracle Database products or options, see the appropriate manual in the product documentation library, available on the Oracle Documentation Library CD-ROM, the DVD-ROM, or on the OTN Web site.

HP-UX Only: Changing Permissions on the extjob Executable

After you run the `root.sh` script during the installation, enter the following commands as `root`:

```
# cd oracle_home/bin
# mv extjob.nobody extjob
# chown extjob extjob
# chmod 4711 extjob
```

In this example, `oracle_home` is the Oracle home directory where you installed the software and `extjob` is the unprivileged user that you created in the "[HP-UX Only: Creating an Unprivileged User](#)" section on page 2-31.

Recommended Post-installation Tasks

Oracle recommends that you perform the tasks described in the following section after completing an installation:

- [Backing Up the root.sh Script](#)
- [Configuring New or Upgraded Databases](#)
- [Setting Up User Accounts](#)
- [Generating the Client Static Library](#)

Backing 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 updates the contents of the existing `root.sh` script during the installation. If you require information contained in the original `root.sh` script, then you can recover it from the backed up `root.sh` file.

Configuring New or Upgraded Databases

Oracle recommends that you run the `utlrp.sql` script after creating or upgrading a database. This script recompiles all PL/SQL modules that might be in an invalid state, including packages, procedures, and types. This is an optional step but Oracle recommends that you do it during installation and not at a later date.

To run the `utlrp.sql` script, follow these steps:

1. Switch user to `oracle`.
2. Use the `oraenv` or `coraenv` script to set the environment for the database where you want to run the `utlrp.sql` script:

- Bourne, Bash, or Korn shell:

```
$ . /usr/local/bin/oraenv
```

- C shell:

```
% source /usr/local/bin/coraenv
```

When prompted, specify the SID for the database.

3. Start SQL*Plus, as follows:

```
$ sqlplus "/ AS SYSDBA"
```

4. If necessary, start the database:

```
SQL> STARTUP
```

5. Run the `utlrp.sql` script:

```
SQL> @?/rdms/admin/utlrp.sql
```

Setting Up User Accounts

For information about setting up additional user accounts, see the *Oracle Database Administrator's Reference for UNIX Systems*.

Generating the Client Static Library

The client static library (`libclntst.a`) is not generated during installation. If you want to link your applications to the client static library, you must first generate it as follows:

1. Switch user to `oracle`.
2. Set the `ORACLE_HOME` environment variable to specify the Oracle home directory used by the Oracle Database installation. For example:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=/u01/app/oracle/product/10.1.0/db_1
$ export ORACLE_HOME
```

- C shell:

```
% setenv ORACLE_HOME /u01/app/oracle/product/10.1.0/db_1
```

3. Enter the following command:

```
$ $ORACLE_HOME/bin/genclntst
```

Required Product-Specific Post-installation Tasks

The following sections describe platform-specific post-installation tasks that you must perform if you installed and intend to use the products mentioned:

- [Configuring Oracle Net Services](#)
- [Configuring Oracle Label Security](#)
- [Installing Natively Compiled Java Libraries for Oracle JVM and Oracle interMedia](#)
- [Installing Oracle Text Supplied Knowledge Bases](#)
- [Configuring Oracle Messaging Gateway](#)
- [Configuring Oracle Precompilers](#)
- [Linux x86 and Linux Itanium Only: Configure OCCI](#)
- [AIX, HP-UX, Solaris SPARC, and zSeries Linux: Configuring the 64-Bit Oracle JDBC/OCI Driver for JDK 1.4](#)

Note: You need only perform post-installation tasks for products that you intend to use.

Configuring Oracle Net Services

If you have a previous release of Oracle software installed on this system, you might want to copy information from the Oracle Net `tnsnames.ora` and `listener.ora` configuration files from the previous release to the corresponding files for the new release.

Note: The default location for the `tnsnames.ora` and `listener.ora` files is the `$ORACLE_HOME/network/admin/` directory. However, you can also use a central location for these files, for example `/var/opt/oracle` on Solaris or `/etc` on other platforms.

Modifying the listener.ora File

If you are upgrading from a previous release of Oracle Database, Oracle recommends that you use the current release of Oracle Net listener instead of the listener from the previous release.

To use the listener from the current release, you may need to copy static service information from the `listener.ora` file from the previous release to the version of that file used by the new release.

For any database instances earlier than release 8.0.3, add static service information to the `listener.ora` file. Oracle Database releases later than release 8.0.3 do not require static service information.

Modifying the tnsnames.ora File

Unless you are using a central `tnsnames.ora` file, copy Oracle Net service names and connect descriptors from the previous release `tnsnames.ora` file to the version of that file used by the new release.

If necessary, you can also add connection information for additional database instances to the new file.

Configuring Oracle Label Security

If you installed Oracle Label Security, you must configure it in a database before you use it. You can configure Oracle Label Security in two ways; with Oracle Internet Directory integration and without Oracle Internet Directory integration. If you configure Oracle Label Security without Oracle Internet Directory integration, you cannot configure it to use Oracle Internet Directory at a later stage.

Note: To configure Oracle Label Security with Oracle Internet Directory integration, Oracle Internet Directory must be installed in your environment and the Oracle database must be registered in the directory.

See Also: For more information about Oracle Label Security enabled with Oracle Internet Directory, see the *Oracle Label Security Administrator's Guide*.

Installing Natively Compiled Java Libraries for Oracle JVM and Oracle *interMedia*

If you plan to use Oracle JVM or Oracle *interMedia*, Oracle strongly recommends that you install the natively compiled Java libraries (NCOMPs) used by those products from the Oracle Database 10g Companion CD. These libraries are required to improve the performance of the products on your platform.

For information about how to install products from the Companion CD, see the "[Installing Oracle Database 10g Products from the Companion CD](#)" section on page 4-11.

Installing Oracle Text Supplied Knowledge Bases

An Oracle Text knowledge base is a hierarchical tree of concepts used for theme indexing, ABOUT queries, and deriving themes for document services. If you plan to use any of these Oracle Text features, you can install two supplied knowledge bases (English and French) from the Oracle Database 10g Companion CD.

Note: You can extend the supplied knowledge bases depending on your requirements. Alternatively, you can create your own knowledge bases, possibly in languages other than English and French. For more information about creating and extending knowledge bases, see the *Oracle Text Reference*.

For information about how to install products from the Companion CD, see the ["Installing Oracle Database 10g Products from the Companion CD"](#) section on page 4-11.

Configuring Oracle Messaging Gateway

Note: Oracle Messaging Gateway is not supported on HP-UX Itanium or Linux Itanium.

To configure Oracle Messaging Gateway, see the section about Messaging Gateway in the *Oracle Streams Advanced Queuing User's Guide and Reference* manual. When following the instructions listed in that manual, refer to this section for additional platform-specific instructions about configuring the `listener.ora`, `tnsnames.ora`, and `mgw.ora` files.

Modifying the listener.ora File for External Procedures

To modify the `$ORACLE_HOME/network/admin/listener.ora` file for external procedures:

1. Back up the `listener.ora` file.
2. Ensure that the default IPC protocol address for external procedures is set as follows:

```
(ADDRESS = (PROTOCOL=IPC)(KEY=EXTPROC))
```
3. Add static service information for a service called `mgwextproc` by adding lines similar to the following to the `SID_LIST` parameter for the listener in the `listener.ora` file:

```
(SID_DESC =  
  (SID_NAME = mgwextproc)  
  (ENVS = platform-specific_env_vars)  
  (ORACLE_HOME = oracle_home)  
  (PROGRAM = extproc_agent)  
)
```


In this example:

- The ENV_S parameter defines the shared library path environment variable and any other required environment variables.

The following table lists the environment variables and required values that you must specify for each platform. In the shared library path environment variable, you must also add any additional library paths required for non-Oracle messaging systems, for example WebSphere MQ or TIBCO Rendezvous.

Platform	ENV_S Parameter Setting
AIX	EXTPROC_DLLS=/oracle_home/lib32/libmgwagent.so, LIBPATH=/oracle_home/jdk/jre/bin:/oracle_home/jdk/jre/bin/classic:/oracle_home/lib32
HP-UX PA-RISC	EXTPROC_DLLS=/oracle_home/lib32/libmgwagent.sl, LD_PRELOAD=/oracle_home/jdk/jre/lib/PA-RISC/server/libjvm.sl, SHLIB_PATH=/oracle_home/jdk/jre/lib/PA_RISC:/oracle_home/jdk/jre/lib/PA_RISC/server: /oracle_home/lib32
Linux x86	EXTPROC_DLLS=/oracle_home/lib/libmgwagent.so, LD_LIBRARY_PATH=/oracle_home/jdk/jre/lib/i386:/oracle_home/jdk/jre/lib/i386/server: /oracle_home/lib
Solaris SPARC	EXTPROC_DLLS=/oracle_home/lib32/libmgwagent.so, LD_LIBRARY_PATH=/oracle_home/jdk/jre/lib/sparc:/oracle_home/lib32
Solaris x86	EXTPROC_DLLS=/oracle_home/lib/libmgwagent.so, LD_LIBRARY_PATH=/oracle_home/jdk/jre/lib/i386:/oracle_home/lib
Tru64 UNIX	EXTPROC_DLLS=/oracle_home/lib/libmgwagent.so, LD_LIBRARY_PATH=/oracle_home/jdk/jre/lib/alpha/classic:/oracle_home/lib
zSeries Linux	EXTPROC_DLLS=/oracle_home/lib32/libmgwagent.so, LD_LIBRARY_PATH=/oracle_home/jdk/jre/bin:/oracle_home/jdk/jre/bin/classic: /oracle_home/lib32

- *oracle_home* is the path of the Oracle home directory.
- *extproc_agent* is the external procedure agent executable file. The following table lists the correct executable file for each platform:

Platform	Agent Executable File
Linux x86, Solaris x86, and Tru64 UNIX	extproc
AIX, HP-UX PA-RISC, Solaris SPARC, and zSeries Linux	extproc32

The following examples show sample `listener.ora` files on Linux x86 and HP-UX PA-RISC:

Note: In the following examples, the `PLSExtProc` service is the default service for PL/SQL external procedures.

- Linux x86:

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME = /u01/app/oracle/product/10.1.0/db_1)
      (PROGRAM = extproc)
    )
    (SID_DESC =
      (SID_NAME = mgwextproc)
      (ENVS = EXTPROC_DLLS=/oracle_home/lib/libmgwagent.so,
        LD_LIBRARY_PATH =/u01/app/oracle/product/10.1.0/db_1/jdk/jre/
        lib/i386:/u01/app/oracle/product/10.1.0/db_1/jdk/jre/lib/i386/server:
        /u01/app/oracle/product/10.1.0/db_1/lib)
      (ORACLE_HOME = /u01/app/oracle/product/10.1.0/db_1)
      (PROGRAM = extproc)
    )
  )
)
```

- HP-UX PA-RISC:

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME = /u01/app/oracle/product/10.1.0/db_1)
      (PROGRAM = extproc)
    )
    (SID_DESC =
      (SID_NAME = mgwextproc)
      (ENVS = EXTPROC_DLLS=/u01/app/oracle/product/10.1.0/db_1/lib32/
        libmgwagent.sl,LD_PRELOAD=/u01/app/oracle/product/10.1.0/db_1/jdk/jre/
        lib/PA-RISC/server/libjvm.sl,SHLIB_PATH=/u01/app/oracle/product/10.1.0/db_1/
        /jdk/jre/lib/PA_RISC:/u01/app/oracle/product/10.1.0/db_1/jdk/jre/lib/PA_RIS
        C/server:/u01/app/oracle/product/10.1.0/db_1/lib32)
      (ORACLE_HOME = /u01/app/oracle/product/10.1.0/db_1)
      (PROGRAM = extproc32)
    )
  )
)
```

Modifying the tnsnames.ora File for External Procedures

To modify the `$ORACLE_HOME/network/admin/tnsnames.ora` file for external procedures:

1. Back up the `tnsnames.ora` file.
2. In the `tnsnames.ora` file, add a connect descriptor with the net service name `MGW_AGENT`, as follows:

```
MGW_AGENT =
(DESCRIPTION=
  (ADDRESS_LIST= (ADDRESS= (PROTOCOL=IPC)(KEY=EXTPROC)))
  (CONNECT_DATA= (SID=mgwextproc) (PRESENTATION=RO)))
```

In this example:

- The value specified for the `KEY` parameter must match the value specified for that parameter in the `IPC` protocol address in the `listener.ora` file.
- The value of the `SID` parameter must match the service name in the `listener.ora` file that you specified for the Oracle Messaging Gateway external procedure agent in the previous section (`mgwextproc`).

Setting up the mgw.ora Initialization File

To modify the `$ORACLE_HOME/mgw/admin/mgw.ora` file for external procedures, set the `CLASSPATH` environment variable to include the classes in the following table and any additional classes required for Oracle Messaging Gateway to access non-Oracle messaging systems, for example WebSphere MQ or TIBCO Rendezvous classes:

Classes	Path
Oracle Messaging Gateway	<code>ORACLE_HOME/mgw/classes/mgw.jar</code>
JRE internationalization	<code>ORACLE_HOME/JRE/lib/i18n.jar</code>
JRE runtime	<code>ORACLE_HOME/JRE/lib/rt.jar</code>
Oracle JDBC	<code>ORACLE_HOME/jdbc/lib/ojdbc14.jar</code>
Oracle internationalization	<code>ORACLE_HOME/jdbc/lib/orai18n.jar</code>
SQLJ	<code>ORACLE_HOME/sqlj/lib/translator.zip</code> <code>ORACLE_HOME/sqlj/lib/runtime12.zip</code>
JMS Interface	<code>ORACLE_HOME/rdbms/jlib/jmscommon.jar</code>
Oracle JMS implementation	<code>ORACLE_HOME/rdbms/jlib/aqapi13.jar</code>
Java Transaction API	<code>ORACLE_HOME/jlib/jta.jar</code>

Configuring Oracle Precompilers

The following section describes post-installation tasks for Oracle precompilers.

Note: All precompiler configuration files are located in the `$ORACLE_HOME/precomp/admin` directory.

Configuring Pro*C/C++

Verify that the PATH environment variable setting includes the directory that contains the C compiler executable. [Table 4–1](#) shows the default directories and the appropriate commands to verify the path setting, depending on your platform and compiler.

Table 4–1 C/C++ Compiler Directory

Platform	Path	Command
AIX and Tru64 UNIX	<code>/usr/bin</code>	<code>\$ which cc</code>
HP-UX	<code>/usr/bin</code>	<code>\$ which cc</code>
Linux x86	<code>/usr/bin</code>	<code>\$ which gcc</code>
	<code>/opt/intel/compiler70/ia32/bin</code>	<code>\$ which icc</code>
Linux Itanium	<code>/usr/bin</code>	<code>\$ which gcc</code>
	<code>/opt/intel/compiler70/ia64/bin</code>	<code>\$ which ecc</code>
Solaris	<code>/opt/SUNWspro/bin</code>	<code>\$ which cc</code>
zSeries Linux	<code>/usr/bin</code>	<code>\$ which gcc</code>

Linux x86 and Linux Itanium Only: Configure OCCI

Two sets of Oracle C++ Call Interface (OCCI) libraries are installed with release 10.1.0.3 on Linux systems. One set is for use with GNU C++ compiler version 2.96 and the other set is for use with GNU C++ version 3.2. The libraries are installed in the following locations:

- Libraries for use with GNU C++ version 2.96:
 - `$ORACLE_HOME/lib/gcc2.96/libocci10.a`
 - `$ORACLE_HOME/lib/gcc2.96/libocci10.so.10.1`
- Libraries for use with GNU C++ version 3.2:
 - `$ORACLE_HOME/lib/gcc3/libocci10.a`
 - `$ORACLE_HOME/lib/gcc3/libocci10.so.10.1`

Although the libraries are installed on all Linux distributions, they are not supported on all distributions. The following table shows the supported GNU C++ version for each distribution:

Linux Distribution	GNU C++ Versions Supported
Red Hat Enterprise Linux AS/ES 2.1	Only 2.96
Red Hat Enterprise Linux AS/ES 3	Both 2.96 and 3.2
SUSE Linux Enterprise Server 8 and 9	Only 3.2

To create the required symbolic links in the `$ORACLE_HOME/lib` directory to the appropriate libraries for your distribution, you must run the following script after all database or client installations:

```
$ORACLE_HOME/bin/occi_links.sh
```

Note: On Red Hat Enterprise Linux 3, this script creates links to the libraries for GNU C++ version 3.2 by default. If you want to use GNU C++ version 2.96, you must manually create the symbolic links to the appropriate libraries. For example:

```
$ ln -s $ORACLE_HOME/lib/gcc2.96/libocci* $ORACLE_HOME/lib
```

AIX, HP-UX, Solaris SPARC, and zSeries Linux: Configuring the 64-Bit Oracle JDBC/OCI Driver for JDK 1.4

To configure the 64-bit Oracle JDBC/OCI Driver for JDK 1.4 on AIX, HP-UX, Solaris SPARC, or zSeries Linux platforms, include the `$ORACLE_HOME/lib` directory in the value specified by the shared library path environment variable for your platform. The following table lists the environment variable for each platform:

Platform	Environment Variable
AIX	LIBPATH
HP-UX and zSeries Linux	LD_LIBRARY_PATH
Solaris	LD_LIBRARY_PATH_64

Installing Oracle Database 10g Products from the Companion CD

The Oracle Database 10g Companion CD contains additional products that you can install. Whether you need to install these products depends on which Oracle Database products or features you plan to use. If you plan to use the following products or features, Oracle strongly recommends that you complete the Oracle Database 10g Products installation from the Companion CD:

- Oracle JVM
- Oracle *interMedia*
- Oracle Text

To install Oracle Database 10g Products from the Companion CD, follow these steps:

Note: For more detailed installation information, see the *Oracle Database Companion CD Installation Guide*, which is available on the Companion CD.

1. Insert the Oracle Database 10g Companion CD or the Oracle Database 10g DVD-ROM into the disc drive.
2. If necessary, log into the system as the user who installed Oracle Database (typically the `oracle` user).

3. To start the Installer, enter the following commands where *directory_path* is the CD-ROM mount point directory or the path of the companion directory on the DVD-ROM:

```
$ cd /tmp
$ /directory_path/runInstaller
```

If the Installer does not appear, see the "[X Windows Display Errors](#)" section on page F-2 for information about troubleshooting.

4. Use the following guidelines to complete the installation:
 - On the Specify File Locations screen, select the Oracle home name and path for the Oracle Database 10g installation where you want to install the products.
 - On the Select a Product to Install screen, select **Oracle Database 10g Products**.
 - If you are prompted to run the `root.sh` script, enter 3 at the prompt that it displays if you do not want to install Legato Single Server Version (not supported on Solaris x86 or zSeries Linux).

Note: For more information about any of the Installer screens, click **Help**.

Reviewing Your Preconfigured Database

This chapter provides information about the default preconfigured database, including information about Oracle database accounts, passwords, and file locations. It includes information about the following topics:

- [Log In to the Oracle Enterprise Manager Database Control](#)
- [Review Accounts and Passwords](#)
- [Unlock and Reset User Passwords](#)
- [Identify Databases](#)
- [Locate the Parameter File](#)
- [Review Tablespaces and Datafiles, Redo Log Files, and Control Files](#)

Log In to the Oracle Enterprise Manager Database Control

If you configured the Oracle Enterprise Manager Database Control during the installation, you can use it to manage your database. Alternatively, you can also use Oracle Enterprise Manager Grid Control to manager your database. To display the Database Control:

1. Use a Web browser to access the Database Control URL:

```
http://host:port/em
```

In this example:

- *host* is the name of the computer on which you installed Oracle Database
- *port* is the port number reserved for the Database Control during installation

If you do not know the correct port number to use, look for the following line in the `$ORACLE_HOME/install/portlist.ini` file:

```
Enterprise Manager Console HTTP Port (db_name) = 5500
```

The installation reserves the first available port from the range 5500 to 5519. For example, if you installed Oracle Database on host `mgmt42`, and the Database Control uses port 5500, enter the following URL:

```
http://mgmt42:5500/em
```

Oracle Enterprise Manager displays the Database Control login page.

2. Log in to the database using the user name `SYS` and connect as `SYSDBA`.

Use the password that you specified for the `SYS` account during the installation.

Note: You can also log in to the Database Control using the `SYSTEM` or `SYSMAN` accounts or you can grant login privileges to other database users.

Understanding Database Control Login Privileges

When you log in to the Oracle Enterprise Manager Database Control using the `SYSMAN` user account, you are logging in as the Oracle Enterprise Manager super user. The `SYSMAN` account is automatically granted the roles and privileges required to access all of the management features provided by the Database Control.

You can also use the `SYS` and `SYSTEM` accounts to log in to the Database Control. In addition, you can grant login privileges to other database users, as follows:

1. Log in to the Database Control.

See Also: See "[Log In to the Oracle Enterprise Manager Database Control](#)" on page 5-2 for information about logging in to the Database Control.

2. Click **Setup** at the top of the Database Control home page.
3. Click **Administrators** in the left navigation bar.
4. Click **Create** to create a new Enterprise Manager user.
5. In the **Name** field, enter the user name of an existing database user or click the flashlight icon and select a user from the pop-up window.

6. Enter the password for this user, then click **Finish**.

Enterprise Manager assigns login privileges to the specified user and includes this user in the list of Enterprise Manager users on the Setup Administrators page.

Review Accounts and Passwords

All databases created by the Database Configuration Assistant (DBCA) include the SYS, SYSTEM, SYSMAN, and DBSNMP database accounts. In addition, Oracle provides several other administrative accounts. Before using these other accounts, you must unlock them and reset their passwords. [Table 5-1](#) describes these accounts, listing their user names and default passwords.

See Also: See the "[Unlock and Reset User Passwords](#)" section on page 5-6 for information about unlocking and resetting passwords.

Note: You can use the Oracle Enterprise Manager Database Control to view the complete list of database accounts. For more information, see "[Unlock and Reset User Passwords](#)" on page 5-6.

Table 5-1 Database Accounts

User Name	Default Password	Description	See Also
ANONYMOUS	ANONYMOUS	Allows HTTP access to Oracle XML DB.	None
BI	BI	The account that owns the Business Intelligence schema included in the Oracle Sample Schemas. It is available only if you loaded the Sample Schemas.	<i>Oracle Database Sample Schemas</i>
CTXSYS	CTXSYS	The Oracle Text account.	<i>Oracle Text Reference</i>
DBSNMP	Password set during installation or in the DBCA	The account used by the Management Agent component of Oracle Enterprise Manager to monitor and manage the database. It is created only if you configure the database to use the Database Control.	<i>Oracle Enterprise Manager Grid Control Installation and Basic Configuration</i>
DIP	DIP	The account used by the Directory Integration Platform (DIP) to synchronize the changes in Oracle Internet Directory with the applications in the database.	<i>Oracle Internet Directory Administrator's Guide</i>
DMSYS	DMSYS	The data mining account. DMSYS performs data mining operations.	<i>Oracle Spatial User's Guide and Reference</i>
EXFSYS	EXFSYS	The account that owns the Expression Filter schema.	None
HR	HR	The account that owns the Human Resources schema included in the Oracle Sample Schemas. It is available only if you loaded the Sample Schemas.	<i>Oracle Database Sample Schemas</i>

Table 5–1 (Cont.) Database Accounts

User Name	Default Password	Description	See Also
IX	IX	The account that owns the Information Transport schema included in the Oracle Sample Schemas. It is available only if you loaded the Sample Schemas.	<i>Oracle Database Sample Schemas</i>
LBACSYS	LBACSYS	The Oracle Label Security administrator account.	<i>Oracle Label Security Administrator's Guide</i>
MDDATA	MDDATA	The schema used by Oracle Spatial for storing Geocoder and router data.	<i>Oracle Spatial User's Guide and Reference</i>
MDSYS	MDSYS	The Oracle Spatial and Oracle <i>interMedia</i> Locator administrator account.	<i>Oracle Spatial User's Guide and Reference</i>
MGMT_VIEW	Randomly generated	An account used by Oracle Enterprise Manager Database Control.	None
ODM	ODM	The data mining account. ODM performs data mining operations.	<i>Oracle Data Mining Administrator's Guide</i> <i>Oracle Data Mining Concepts</i>
ODM_MTR	MTRPW	The account associated with the data repository for data mining sample programs.	<i>Oracle Data Mining Administrator's Guide</i> <i>Oracle Data Mining Concepts</i>
OE	OE	The account that owns the Order Entry schema included in the Oracle Sample Schemas. It is available only if you loaded the Sample Schemas.	<i>Oracle Database Sample Schemas</i>
OLAPSYS	MANAGER	The account that owns the OLAP catalogs.	<i>Oracle OLAP Application Developer's Guide</i>
ORDPLUGINS	ORDPLUGINS	The Oracle <i>interMedia</i> user. Plugins supplied by Oracle and third-party plugins are installed in this schema.	<i>Oracle interMedia Reference</i>
ORDSYS	ORDSYS	The Oracle <i>interMedia</i> administrator account.	<i>Oracle interMedia Reference</i>
OUTLN	OUTLN	The account that supports plan stability. Plan stability allows you to maintain the same execution plans for the same SQL statements. OUTLN acts as a role to centrally manage metadata associated with stored outlines.	<i>Oracle Database Concepts</i>

Table 5–1 (Cont.) Database Accounts

User Name	Default Password	Description	See Also
PM	PM	The account that owns the Product Media schema included in the Oracle Sample Schemas. It is available only if you loaded the Sample Schemas.	<i>Oracle Database Sample Schemas</i>
SCOTT	TIGER	An account used by Oracle sample programs and examples.	<i>Oracle Database Administrator's Guide</i>
SH	SH	The account that owns the Sales History schema included in the Oracle Sample Schemas. It is available only if you loaded the Sample Schemas during an Enterprise Edition installation.	<i>Oracle Database Administrator's Guide</i>
SI_INFORMTN_SCHEMA	SI_INFORMTN_SCHEMA	The account that stores the information views for the SQL/MM Still Image Standard.	<i>Oracle interMedia Reference</i>
SYS	Password set during installation or in the DBCA	The account used to perform database administration tasks.	<i>Oracle Database Administrator's Guide</i>
SYSMAN	Password set during installation or in the DBCA	The account used to perform Oracle Enterprise Manager database administration tasks. It is created only if you configure the database to use the Database Control.	<i>Oracle Enterprise Manager Grid Control Installation and Basic Configuration</i>
SYSTEM	Password set during installation or in the DBCA	Another account used to perform database administration tasks.	<i>Oracle Database Administrator's Guide</i>
WMSYS	WMSYS	The account used to store the metadata information for Oracle Workspace Manager.	<i>Oracle Database Application Developer's Guide - Workspace Manager</i>
WKPROXY	WKSYS	The Ultra Search proxy user.	<i>Oracle Ultra Search User's Guide</i>
WK_TEST	WK_TEST	The default Ultra Search instance schema.	<i>Oracle Ultra Search User's Guide</i>
WKSYS	WMSYS	The account used to store Ultra Search system dictionaries and PL/SQL packages.	<i>Oracle Ultra Search User's Guide</i>
XDB	XDB	The account used for storing Oracle XML DB data and metadata.	<i>Oracle XML DB Developer's Guide</i>

Unlock and Reset User Passwords

Passwords for all Oracle system administration accounts except SYS, SYSTEM, SYSMAN, and DBSMP are revoked after installation. Before you use a locked account, you must unlock it and reset its password. If you created a preconfigured database during the installation, but you did not unlock a required account, you must unlock it now, using one of the following methods:

- Using Database Control to Unlock Accounts and Reset Passwords
- Using SQL*Plus to Unlock Accounts and Reset Passwords

Note: If you are creating a new database using DBCA, you can unlock accounts after the database is created by clicking **Password Management** before you exit from DBCA.

Using Database Control to Unlock Accounts and Reset Passwords

To unlock and reset user account passwords using Oracle Enterprise Manager Database Control:

1. Log in to the Database Control.

See Also: See "[Log In to the Oracle Enterprise Manager Database Control](#)" on page 5-2 for information about logging in to the Database Control.

2. Click **Administration**.

3. In the Security section of the Administration page, click **Users**.

Enterprise Manager displays a table listing all database accounts. The Account Status column indicates whether the account is locked and whether the password is expired.

4. Select the user account that you want to modify, then click **Edit**.
5. Use the General page of the Users property sheet to unlock the account and, optionally, to change the password.

See Also: Click **Help** in the Database Control window for more information about using the Database Control.

Using SQL*Plus to Unlock Accounts and Reset Passwords

To unlock and reset user account passwords using SQL*Plus:

1. Start SQL*Plus and log in as the SYS user, connecting as SYSDBA:

```
$ sqlplus /nolog
SQL> CONNECT SYS/SYS_password AS SYSDBA
```

2. Enter a command similar to the following, where *account* is the user account that you want to unlock and *password* is the new password:

```
SQL> ALTER USER account [ IDENTIFIED BY password ] ACCOUNT UNLOCK;
```

In this example:

- The ACCOUNT UNLOCK clause unlocks the account.
- The IDENTIFIED BY *password* clause resets the password.

Note: If you unlock an account but do not reset the password, then the password remains expired. The first time someone connects as that user, they must change the user's password.

To permit unauthenticated access to your data through HTTP, unlock the ANONYMOUS user account.

See Also: See the *Oracle Database Administrator's Guide* for more information about:

- Unlocking and changing passwords after installation
- Oracle security procedures
- Best security practices

Identify Databases

The Oracle Database software identifies a database by its global database name. A global database name consists of the database name and database domain. Usually, the database domain is the same as the network domain, but it need not be. The global database name uniquely distinguishes a database from any other database in the same network. You specify the global database name when you create a database during the installation, or using the Database Configuration Assistant. For example:

```
sales.us.oracle.com
```

In this example:

- `sales` is the name of the database. The database name portion is a string of no more than 8 characters that can contain alphanumeric, underscore (`_`), and pound (`#`) characters. The `DB_NAME` initialization parameter specifies the database name.
- `us.acme.com` is the database domain in which the database is located. In this example, the database domain is the same as the network domain. Together, the database name and the database domain make the global database name unique. The domain portion is a string of no more than 128 characters that can contain alphanumeric, underscore (`_`), and pound (`#`) characters. The `DB_DOMAIN` initialization parameter specifies the database domain name.

The `DB_NAME` parameter and the `DB_DOMAIN` name parameter combine to create the global database name value assigned to the `SERVICE_NAMES` parameter in the initialization parameter file.

The System Identifier (SID) identifies a specific database instance. The SID uniquely distinguishes the instance from any other instance on the same computer. Each database instance requires a unique SID and database name. In most cases, the SID is the same as the database name portion of the global database name.

Locate the Parameter File

By default, the preconfigured database uses a server parameter file named `spfilesid.ora`, which is stored in the `$ORACLE_HOME/dbs` directory. However, if you choose ASM or raw device storage for the database, DBCA typically uses the same storage mechanism for the server parameter file.

If the server parameter file is not located in the `$ORACLE_HOME/dbs` directory, the database uses the `SPFILE` parameter in an initialization parameter file to locate it. The default initialization parameter file is `$ORACLE_HOME/dbs/initsid.ora`.

You can use the Oracle Enterprise Manager Database Control to view the location of the server parameter file and list all of the initialization parameters, as follows:

1. Log in to the Database Control.

See Also: See "[Log In to the Oracle Enterprise Manager Database Control](#)" on page 5-2 for information about logging in to the Database Control.

2. Click **Administration**.
3. In the Instance section of the Administration page, click **All Initialization Parameters**.

Enterprise Manager displays a table listing the current value of each initialization parameter.

4. Click **SPFile**.

Enterprise Manager displays a table listing the value of each initialization parameter specified in the server parameter file. The location of the server parameter file is displayed above the table.

Review Tablespaces and Datafiles, Redo Log Files, and Control Files

The following sections contain information about tablespaces and datafiles, redo log files, and control files:

- [Identifying Tablespaces and Datafiles](#)
- [Locating Redo Log Files](#)
- [Locating Control Files](#)

Identifying Tablespaces and Datafiles

An Oracle database is divided into smaller logical areas of space known as tablespaces. Each tablespace corresponds to one or more physical datafiles. Datafiles contain the contents of logical database structures such as tables and indexes. You can associate each datafile with only one tablespace and database.

Note: The SYSAUX and SYSTEM tablespaces must be present in all Oracle Database 10g databases.

[Table 5–2](#) describes the tablespaces provided by the default preconfigured database.

Table 5–2 Tablespaces and Descriptions

Tablespace	Description
EXAMPLE	Stores the sample schemas, if you included them.
SYSAUX	Serves as an auxiliary tablespace to the SYSTEM tablespace. Some products and options that previously used the SYSTEM tablespace now use the SYSAUX tablespace, which reduces the load on the SYSTEM tablespace.
SYSTEM	Stores the data dictionary, which includes definitions of tables, views, and stored procedures needed by Oracle Database.
TEMP	Stores temporary tables and indexes when processing SQL statements.
UNDOTBS1	Stores undo information.
USERS	Stores database objects created by database users.

See Also: See *Oracle Database Concepts* and the *Oracle Database Administrator's Guide* for more information about tablespaces and datafiles.

To use the Oracle Enterprise Manager Database Control to view the list of datafiles used by the database and their associated tablespaces:

1. Log in to the Database Control.

See Also: See "[Log In to the Oracle Enterprise Manager Database Control](#)" on page 5-2 for information about logging in to the Database Control.

2. Click **Administration**.
3. In the Storage section of the Administration page, click **Datafiles**.

Enterprise Manager displays a table listing each datafile, and the tablespace with which it is associated.

See Also: For more information about using the Database Control to view, modify, and create tablespaces, click **Help** in the Database Control window.

Locating Redo Log Files

The preconfigured database uses three redo log files. Redo log files record all changes made to data in the database buffer cache. If an instance fails, then Oracle Database uses the redo log files to recover the modified data in memory.

Oracle Database uses redo log files in a cyclical fashion. For example, if three files constitute the online redo log, Oracle Database fills the first file, then the second file, and then the third file. In the next cycle, it reuses and fills the first file, the second file, and so on.

See Also: See *Oracle Database Backup and Recovery Basics* for more information about redo log files.

To use the Oracle Enterprise Manager Database Control to view or modify the redo log files for your preconfigured database:

1. Log in to the Database Control.

See Also: See "[Log In to the Oracle Enterprise Manager Database Control](#)" on page 5-2 for information about logging in to the Database Control.

2. Click **Administration**.
3. In the Storage section of the Administration page, click **Redo Log Groups**.
Enterprise Manager displays a table listing the redo log groups used by the database.
4. To view the name and location of the redo log file associated with a particular group, select that group then click **View**.

See Also: For more information about using the Database Control to view, modify, and create redo log files, click **Help** in the Database Control window.

Locating Control Files

The preconfigured database uses three control files. Oracle recommends that you keep at least three control files for each database and set the CONTROL_FILES initialization parameter to specify the location of each file.

A control file is an administrative file. Oracle Database requires a control file to start and run the database. The control file defines the physical structure of the database. For example, it defines the database name and the names and locations of the database datafiles and redo log files.

To use the Oracle Enterprise Manager Database Control to view information about the control files for your preconfigured database:

1. Log in to the Database Control.

See Also: See "[Log In to the Oracle Enterprise Manager Database Control](#)" on page 5-2 for information about logging in to the Database Control.

2. Click **Administration**.
3. In the Storage section of the Administration page, click **Controlfiles**.

Enterprise Manager displays a table listing the control files used by the database.

See Also: For more information about using the Database Control to view information about control files and back them up to trace files, click **Help** in the Database Control window.

For more information about setting the CONTROL_FILES initialization parameter value, see the *Oracle Database Administrator's Guide*.

Removing Oracle Software

This chapter describes how to completely remove all Oracle databases, instances, and software from an Oracle home directory. It includes information about the following topics:

- [Overview](#)
- [Identifying All Instances](#)
- [Removing an Oracle Database](#)
- [Removing an ASM Instance](#)
- [Reconfiguring Oracle Cluster Synchronization Services](#)
- [Removing Oracle Software](#)

Note: If you want to remove an Oracle Real Application Clusters (RAC) installation, see the *Oracle Real Application Clusters Installation and Configuration Guide* for more information.

If you want to remove an individual product, check the product-specific documentation for requirements and restrictions.

Overview

To completely remove all Oracle databases, instances, and software from an Oracle home directory, you must:

- Identify all instances associated with the Oracle home.
- Remove database and ASM instances.
- Shut down processes.
- Reconfigure the Oracle Cluster Synchronization Services Daemon, if necessary.
- Remove the Oracle software.

The following sections describe how to perform these tasks.

Identifying All Instances

To identify all instances associated with the Oracle home that you want to remove, enter one of the following commands:

- Solaris:

```
$ more /var/opt/oracle/oratab
```
- Other operating systems:

```
$ more /etc/oratab
```

The output for this command contains entries similar to the following:

```
+ASM:/u01/app/oracle/product/10.1.0/db_1:N  
CUST:/u01/app/oracle/product/10.1.0/db_1:N
```

These entries show that the +ASM Automatic Storage Manager instance and the CUST Oracle database instance are associated with the /u01/app/oracle/product/10.1.0/db_1 Oracle home directory.

Removing an Oracle Database

To completely remove Oracle Database software, you must remove any installed databases. To remove an Oracle database:

Note: Removing an Oracle database deletes all of the data in the database. If you want to keep this data, make sure that you back up the database before deleting it.

1. Log in as the `oracle` user:

```
$ su - oracle
```
2. Run the `oraenv` or `coraenv` script to set the environment for the database that you want to remove, for example:
 - Bourne, Bash, or Korn shell:

```
$ . /usr/local/bin/oraenv
```
 - C shell:

```
% source /usr/local/bin/coraenv
```

3. At the prompt, specify the SID for the database that you want to remove.
4. Start the Database Configuration Assistant (DBCA):


```
$ dbca
```

The Welcome window appears.
5. Click **Next**.

The Operations window appears.
6. Select **Delete a Database**, then click **Next**.
7. Select the database that you want to delete, then click **Finish**.
8. In the window that appears, confirm that you want to delete the database.
9. When the DBCA removes the database, you are prompted to choose whether you want to perform another operation. Click **Yes** to return to the Operations screen or click **No** to exit from DBCA. If you want to remove another database, click **Yes** and repeat steps 6 to 8.

Removing an ASM Instance

To completely remove Oracle database software, you must also remove any ASM instances running in the Oracle home. To remove an ASM instance:

1. If necessary, log in as the `oracle` user:


```
$ su - oracle
```
2. Run the `oraenv` or `coraenv` script to set the environment for the ASM instance that you want to remove, for example:
 - Bourne, Bash, or Korn shell:


```
$ . /usr/local/bin/oraenv
```
 - C shell:


```
$ source /usr/local/bin/coraenv
```
3. At the prompt, specify the SID for the ASM instance that you want to remove.
4. Connect to the ASM instance as the SYS user:


```
$ sqlplus "SYS/SYS_PASSWORD AS SYSDBA"
```
5. Enter the following command to determine whether any Oracle database instance is using the ASM instance:


```
SQL> SELECT INSTANCE_NAME FROM V$ASM_CLIENT;
```

This command lists all of the database instances that are using this ASM instance.

Note: This command only lists 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, do not remove the ASM instance or the Oracle home.

6. If there are no database instances associated with this ASM instance, drop the disk groups associated with this instance as follows:

Note: Dropping an ASM disk group makes the disk devices that it used available for use with another ASM instance if required. However, all data in the disk group is lost. Make sure that no other database instance requires any data from this disk group before you drop it.

- a. Identify the disk groups associated with the ASM instance:

```
SQL> SELECT NAME FROM V$ASM_DISKGROUP;
```

- b. For each disk group that you want to delete, enter a command similar to the following:

```
SQL> DROP DISKGROUP name INCLUDING CONTENTS;
```

7. Enter the following command to shut down the ASM instance:

```
SQL> SHUTDOWN
```

8. Remove the entry for the ASM instance from the `/etc/oratab` file (or from the `/var/opt/oracle/oratab` file on Solaris).

Reconfiguring Oracle Cluster Synchronization Services

Oracle Cluster Synchronization Services (CSS) is a daemon process that is configured by the `root.sh` script when you install Oracle Database 10g for the first time. It is configured to start every time the system boots. This daemon process is required to enable synchronization between Oracle ASM and database instances. It must be running if an Oracle database is using ASM for database file storage.

Note: On cluster systems with Oracle Real Application Clusters (RAC) installations, the CSS daemon is configured during the Oracle Cluster Ready Services (CRS) installation. If the system is running Oracle CRS, see the *Oracle Real Application Clusters Installation and Configuration Guide* for information about removing RAC or CRS.

Before you remove an Oracle Database 10g Oracle home, you must determine whether the CSS daemon is running from that Oracle home and whether any other Oracle Database 10g Oracle homes exist on the system:

- If the Oracle Database 10g Oracle home that you want to remove is the only Oracle Database 10g installation on the system, you can delete the CSS daemon configuration.
- If the CSS daemon is running from the Oracle Database 10g Oracle home that you want to remove and other Oracle Database 10g installations exist on the system, you must reconfigure the CSS daemon to run from another Oracle Database 10g Oracle home.

The following subsections describe how to complete these tasks.

Identifying Oracle Database 10g Oracle Homes

To identify all of the Oracle Database 10g Oracle home directories, enter one of the following commands:

- Solaris:


```
$ more /var/opt/oracle/oratab
```
- Other operating systems:


```
$ more /etc/oratab
```

From the output, identify any Oracle home directories where Oracle Database 10g is installed. Oracle homes that contain Oracle Database 10g typically have paths similar to the following. However, they might use different paths.

```
/mount_point/app/oracle/product/10.1.0/db_n
```

If there is only one Oracle home directory that contains Oracle Database 10g, see the ["Deleting the Oracle CSS Daemon Configuration"](#) section on page 6-6 for information about deleting the Oracle CSS daemon configuration.

If you identify more than one Oracle Database 10g Oracle home directory, see the following section for information about reconfiguring the Oracle CSS daemon.

Reconfiguring the Oracle CSS Daemon

To reconfigure the Oracle CSS daemon so that it runs from an Oracle home that you are not removing, follow these steps:

1. In all Oracle home directories on the system, stop all Oracle ASM instances and any Oracle Database instances that use ASM for database file storage.
2. Switch user to root.
3. Depending on your operating system, enter one of the following commands to identify the Oracle home directory being used to run the CSS daemon:
 - AIX or Linux:


```
# more /etc/oracle/ocr.loc
```
 - Other operating systems:


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

The output from this command is similar to the following:

```
ocrconfig_loc=/u01/app/oracle/product/10.1.0/db_1/cdata/localhost/local.ocr
local_only=TRUE
```

The `ocrconfig_loc` parameter specifies the location of the Oracle Cluster Registry (OCR) used by the CSS daemon. The path up to the `cdata` directory is the Oracle home directory where the CSS daemon is running (`/u01/app/oracle/product/10.1.0/db_1` in this example).

Note: If the value for the `local_only` parameter is `FALSE`, Oracle CRS is installed on this system. See the *Oracle Real Application Clusters Installation and Configuration Guide* for information about removing RAC or CRS.

If this Oracle home directory is not the Oracle home that you want to remove, you can continue to the "[Removing Oracle Software](#)" section on page 6-7.

4. Change directory to the Oracle home directory for an Oracle Database 10g installation that you are *not* removing.
5. Set the `ORACLE_HOME` environment variable to specify the path to this Oracle home directory:

- Bourne, Bash, or Korn shell:

```
# ORACLE_HOME=/u01/app/oracle/product/10.1.0/db_2
# export ORACLE_HOME
```

- C shell:

```
# setenv ORACLE_HOME /u01/app/oracle/product/10.1.0/db_2
```

6. Enter the following command to reconfigure the CSS daemon to run from this Oracle home:

```
# $ORACLE_HOME/bin/localconfig reset $ORACLE_HOME
```

The script stops the Oracle CSS daemon, reconfigures it in the new Oracle home, and then restarts it. When the system boots, the CSS daemon starts automatically from the new Oracle home.

7. To remove the original Oracle home directory, see the "[Removing Oracle Software](#)" section on page 6-7.

Deleting the Oracle CSS Daemon Configuration

To delete the Oracle CSS daemon configuration, follow these steps:

Note: Delete the CSS daemon configuration only if you are certain that no other Oracle Database 10g installation requires it.

1. Remove any databases or ASM instances associated with this Oracle home. See the preceding sections for information about how to complete these tasks.
2. Switch user to `root`.
3. Change directory to the Oracle home directory that you are removing.

4. Set the ORACLE_HOME environment variable to specify the path to this Oracle home directory:

- Bourne, Bash, or Korn shell:

```
# ORACLE_HOME=/u01/app/oracle/product/10.1.0/db_1
# export ORACLE_HOME
```

- C shell:

```
# setenv ORACLE_HOME /u01/app/oracle/product/10.1.0/db_1
```

5. Enter the following command to delete the CSS daemon configuration from this Oracle home:

```
# $ORACLE_HOME/bin/localconfig delete
```

The script stops the Oracle CSS daemon, then deletes its configuration. When the system boots, the CSS daemon no longer starts.

Removing Oracle Software

The following steps describe how to use the Installer to remove Oracle software from an Oracle home:

Note: Always use the Oracle Universal Installer to remove Oracle software. Do not delete any Oracle home directories without first using the Installer to remove the software.

1. If necessary, log in as the `oracle` user:

```
$ su - oracle
```

2. Set the ORACLE_HOME environment variable to specify the path of the Oracle home directory that you want to remove:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_HOME=/u01/app/oracle/product/10.1.0/db_1
$ export ORACLE_HOME
```

- C shell:

```
$ setenv ORACLE_HOME /u01/app/oracle/product/10.1.0/db_1
```

3. Remove any databases or ASM instances associated with this Oracle home and delete or reconfigure the Oracle CSS daemon.

See the preceding sections for information about how to complete these tasks.

4. Stop any processes running in this Oracle home:

Process Name	Command
Database Control	<code>\$ORACLE_HOME/bin/emctl stop dbconsole</code>
Oracle Net listener	<code>\$ORACLE_HOME/bin/lsnrctl stop</code>
iSQL*Plus	<code>\$ORACLE_HOME/bin/isqlplusctl stop</code>
Ultra Search	<code>\$ORACLE_HOME/bin/searchctl stop</code>

5. Start the Installer as follows:

```
$ $ORACLE_HOME/oui/bin/runInstaller
```

6. In the Welcome window, click **Deinstall Products**.

The Inventory window appears, listing all of the Oracle homes on the system.

7. In the Inventory window, select the Oracle home and the products that you want to remove, then click **Remove**.

Note: If you choose to remove Oracle JVM, the Oracle Universal Installer removes all installed products that depend on Oracle JVM, including Oracle Database.

8. When the Installer displays a confirmation window asking you to confirm that you want to deinstall the products and their dependant components, click **Yes**.

The Installer displays a progress indicator as it removes the software.

9. When the products have been deleted, click **Cancel** to exit from the Installer, then click **Yes**.

Mounting Discs

This Oracle product is delivered on both CD-ROM and DVD-ROM. These discs are in ISO 9660 format with Rockridge extensions. This appendix describes how to mount and unmount a disc on each supported platform. It includes information about the following topics:

- [Mounting Discs on AIX](#) on page A-2
- [Mounting Discs on HP-UX](#) on page A-2
- [Mounting Discs on Linux](#) on page A-3
- [Mounting Discs on Solaris](#) on page A-4
- [Mounting Discs on Tru64 UNIX](#) on page A-5

Mounting Discs on AIX

To mount a disc, follow these steps:

1. Switch user to root:

```
$ su - root
```

2. If necessary, enter a command similar to following to unmount the currently mounted disc, then remove it from the drive:

```
# umount /cdrom
```

In this example, /cdrom is the mount point directory for the disc drive.

3. Insert the appropriate disc into the disc drive, then enter a command similar to the following to mount it:

```
# /usr/sbin/mount -rv cdrfs /dev/cd0 /cdrom
```

In this example, /dev/cd0 is the device name of the disc drive and /cdrom is the mount point directory.

4. If the Installer is displaying the Disk Location dialog box, enter the disc mount point directory path, for example:

```
/cdrom
```

To continue, go to one of the following sections:

- If you want to copy software to a hard disk, see the "[Copy the Oracle Database Software to a Hard Disk](#)" section on page 3-3.
- If you want to install the software from the discs, see the "[AIX Only: Run the rootpre.sh Script](#)" section on page 3-5.

Mounting Discs on HP-UX

To mount a disc, follow these steps:

1. Switch user to root:

```
$ su - root
```

2. If necessary, enter the following command to unmount the currently mounted disc, then remove it from the drive:

```
# /usr/sbin/umount /SD_CDRROM
```

In this example, /SD_CDRROM is the mount point directory for the disc drive.

3. Insert the appropriate disc into the disc drive, then enter a command similar to the following to mount it:

```
# /usr/sbin/mount -F cdfs -o rr /dev/dsk/cxytdz /SD_CDRROM
```

In this example, /SD_CDRROM is the disc mount point directory and /dev/dsk/cxytdz is the device name for the disc device, for example /dev/dsk/c0t2d0.

4. If the Installer is displaying the Disk Location dialog box, enter the disc mount point directory path, for example:

```
/SD_CDRROM
```

To continue, go to one of the following sections:

- If you want to copy software to a hard disk, see the ["Copy the Oracle Database Software to a Hard Disk"](#) section on page 3-3.
- If you want to install the software from the discs, see the ["Install the Oracle Database Software"](#) section on page 3-5.

Mounting Discs on Linux

On most Linux systems, the disc mounts automatically when you insert it into the disc drive. If the disc does not mount automatically, follow these steps to mount it:

1. Switch user to root:

```
$ su - root
```

2. If necessary, enter a command similar to one of the following to eject the currently mounted disc, then remove it from the drive:

- Red Hat:

```
# eject /mnt/cdrom
```

- SUSE:

```
# eject /media/cdrom
```

In these examples, `/mnt/cdrom` and `/media/cdrom` are the mount point directories for the disc drive.

3. Insert the appropriate disc into the disc drive.
4. To verify that the disc mounted automatically, enter one of the following commands depending on your platform:

- Red Hat:

```
# ls /mnt/cdrom
```

- SUSE:

```
# ls /media/cdrom
```

5. If this command fails to display the contents of the disc, enter a command similar to the following to mount it, depending on your platform:

- Red Hat:

```
# mount -t iso9660 /dev/cdrom /mnt/cdrom
```

- SUSE:

```
# mount -t iso9660 /dev/cdrom /media/cdrom
```

In these examples, `/mnt/cdrom` and `/media/cdrom` are the mount point directories for the disc drive.

6. If the Installer is displaying the Disk Location dialog box, enter the disc mount point directory path, for example:

```
/mnt/cdrom
```

To continue, go to one of the following sections:

- If you want to copy software to a hard disk, see the "[Copy the Oracle Database Software to a Hard Disk](#)" section on page 3-3.
- If you want to install the software from the discs, see the "[Install the Oracle Database Software](#)" section on page 3-5.

Mounting Discs on Solaris

On most Solaris systems, the disc mounts automatically when you insert it into the disc drive. If the disc does not mount automatically, follow these steps to mount it:

1. Switch user to `root`:

```
$ su - root
```

2. If necessary, enter the following command to eject the currently mounted disc, then remove it from the drive:

```
# eject
```

3. Insert the appropriate disc into the disc drive.

4. To verify that the disc mounted automatically, enter the following command:

```
# ls /cdrom/cdrom0
```

5. If this command fails to display the contents of the disc, enter a command similar to the following to mount it:

```
# /usr/sbin/mount -r -F hsfs /dev/dsk/cxytdzs2 /cdrom
```

In this example, `/cdrom` is the disc mount point directory and `/dev/dsk/cxytdzs2` is the device name for the disc device, for example `/dev/dsk/c0t2d0s2`.

6. If the Installer is displaying the Disk Location dialog box, enter the disc mount point directory path, for example:

- Disc mounted automatically:

```
/cdrom/cdrom0
```

- Disc mounted manually:

```
/cdrom
```

To continue, go to one of the following sections:

- If you want to copy software to a hard disk, see the "[Copy the Oracle Database Software to a Hard Disk](#)" section on page 3-3.
- If you want to install the software from the discs, see the "[Install the Oracle Database Software](#)" section on page 3-5.

Mounting Discs on Tru64 UNIX

To mount a disc, follow these steps:

1. Switch user to root:

```
$ su - root
```

2. If necessary, enter a command similar to the following to unmount the currently mounted disc, then remove it from the drive:

```
# /usr/sbin/umount /cdrom
```

In this example, /cdrom is the mount point directory for the disc drive.

3. Insert the appropriate disc into the disc drive, then enter a command similar to the following to mount it:

```
# /usr/sbin/mount -t cdfs -o nodefperm,noversion /dev/disk/cdrom0c /cdrom
```

In this example, /cdrom is the disc mount point directory, which must exist, and /dev/disk/cdrom0c is the disc device name.

4. If the Installer is displaying the Disk Location dialog box, enter the disc mount point directory path, for example:

```
/cdrom
```

To continue, go to one of the following sections:

- If you want to copy software to a hard disk, see the "[Copy the Oracle Database Software to a Hard Disk](#)" section on page 3-3.
- If you want to install the software from the discs, see the "[Install the Oracle Database Software](#)" section on page 3-5.

Installing and Configuring Oracle Products Using Response Files

This chapter describes how to install and configure Oracle products using response files. It includes information about the following topics:

- [Introduction](#)
- [Create the oraInst.loc File](#)
- [Prepare a Response File](#)
- [Run the Installer in Silent or Suppressed Mode](#)
- [Run NetCA in Silent Mode](#)
- [Run DBCA in Non-interactive or Silent Mode](#)

Introduction

You can automate the installation and configuration of Oracle software, either fully or partially, by specifying a response file when you start the Oracle Universal Installer. The Installer uses the values contained in the response file to provide answers to some or all of the Installer prompts:

- If you include responses for all of the prompts in the response file and specify the `-silent` option when starting the Installer, then the Installer runs in silent mode. During a silent-mode installation, the Installer does not display any screens. Instead, it displays progress information in the terminal that you used to start it.
- If you include responses for some or all of the prompts in the response file and omit the `-silent` option, then the Installer runs in suppressed mode. During a suppressed-mode installation, the Installer displays only the screens for which you did not specify all required information. You can also use variables in the response file or command-line options to suppress other installer screens, such as the Welcome screen or Summary screen, that do not prompt for information.

The following table describes several reasons why you might want to run the Installer in silent mode or suppressed mode:

Mode	Uses
Silent	<p>Use silent mode if you want to:</p> <ul style="list-style-type: none"> ■ Complete an unattended installation, which you might schedule using operating system utilities such as <code>at</code> ■ Complete several similar installations on multiple systems without user interaction ■ Install the software on a system that does not have X Window System software installed on it <p>The Installer displays progress information in the terminal that you used to start it, but it does not display any of the Installer screens.</p>
Suppressed	<p>Use suppressed mode if you want to complete similar Oracle software installations on more than one system, providing default answers to some, but not all of the Installer prompts.</p> <p>If you do not specify information required for a particular Installer screen in the response file, the Installer displays that screen. It suppresses screens for which you have provided all of the required information.</p>

Creating a Database Using ASM as the Storage Option for Database Files

Before you create a database that uses ASM, you must run the `root.sh` script. For this reason, you cannot create a database using ASM as the storage option for database files during a silent-mode installation. Instead, you can complete a software-only installation using silent-mode, then run the NetCA and DBCA configuration assistants in silent mode after you have completed the software-only installation and you have run the `root.sh` script.

Note: This limitation applies only to databases that use ASM as the storage option for database files. You can create a database that uses the file system or raw devices storage option during a silent-mode installation.

Installation Overview

To install and configure Oracle products using the Installer in silent or suppressed mode, follow these steps:

Note: You must complete all required pre-installation tasks on a system before running the Installer in silent or suppressed mode.

1. Create the oraInst.loc file.
2. Prepare a response file.
3. Run the Installer in silent or suppressed mode.
4. If you completed a software-only installation, run NetCA and DBCA in silent or non-interactive mode if required.

These steps are described in the following sections.

Create the oraInst.loc File

If you plan to install Oracle products using the Installer in silent or suppressed mode, you must manually create the oraInst.loc file if it does not already exist. This file specifies the location of the Oracle Inventory directory where the Installer creates the inventory of Oracle products installed on the system.

Note: If Oracle software has been installed previously on the system, the oraInst.loc file might already exist. If the file does exist, you do not need to create a new file.

To create the oraInst.loc file, follow these steps:

1. Switch user to root:


```
$ su - root
```
2. On HP-UX, Solaris, Tru64 UNIX, and zSeries Linux, create the /var/opt/oracle directory if it does not exist:


```
# mkdir -p /var/opt/oracle
```
3. Change directory as follows, depending on your operating system:
 - AIX, Linux x86, or Linux Itanium:


```
# cd /etc
```
 - Other operating systems:


```
# cd /var/opt/oracle
```
4. Use a text editor to create the oraInst.loc file, containing the following lines:


```
inventory_loc=ORACLE_BASE/oraInventory
inst_group=
```

In this example, *ORACLE_BASE* is the path of the Oracle base directory, for example, /01/app/oracle.

5. Enter the following commands to set the appropriate owner, group, and permissions on the `oraInst.loc` file:

```
# chown oracle:oinstall oraInst.loc
# chmod 664 oraInst.loc
```

Prepare a Response File

There are two methods that you can use to prepare a response file for use during silent-mode or suppressed-mode installations:

- Edit a response file template

This method is most useful for the Enterprise Edition or Standard Edition installation types.
- Record a response file using the Installer

This method is most useful for Custom or software-only installations.

The following subsections describe each method.

Editing a Response File Template

Oracle provides response file templates for each product and installation type, and for each configuration tool. These files are located in the `response` directory on the CD-ROM or the `db/response` directory on the DVD-ROM.

Note: If you copied the software to a hard disk, the response files are located in the `Disk1/response` directory.

Table B–1 lists the response files provided with Oracle Database.

Table B–1 Response Files

Response File	Description
<code>enterprise.rsp</code>	Enterprise Edition installation of Oracle Database
<code>standard.rsp</code>	Standard Edition installation of Oracle Database
<code>custom.rsp</code>	Custom installation of Oracle Database

To prepare a response file:

1. Copy the response file from the response file directory to a directory on your system:

```
$ cp /directory_path/response/response_file.rsp local_directory
```

In this example, `directory_path` is the CD-ROM mount point directory or the `db` directory on the DVD-ROM. If you have copied the software to a hard drive, you can edit the file in the `response` directory if you prefer.

2. Open the response file in a text editor:

```
$ vi /local_dir/response_file.rsp
```

3. Edit the file, following the instructions in the file.

Note: The Installer or configuration assistant fails if you do not correctly configure the response file. See "[Silent-Mode Response File Error Handling](#)" on page F-4 for more information about troubleshooting a failed silent-mode installation.

4. Change the permissions on the file to 700:

```
$ chmod 700 /local_dir/response_file.rsp
```

Caution: A fully specified response file for an Oracle Database installation contains the passwords for database administrative accounts and for a user that is a member of the OSDBA group (required for automated backups). Make sure that only the Oracle software owner user can view or modify response files or consider deleting them after the installation succeeds.

Recording a Response File

You can use the Installer in interactive mode to record a response file that you can edit and then use to complete silent-mode or suppressed-mode installations. When you are recording the response file, you can either complete the installation, or you can exit from the Installer on the Summary page, before it starts to copy the software to the system.

Note: You cannot use record mode to create a response file during an installation that uses the Basic Installation method.

To record a new response file:

1. Complete the pre-installation tasks listed in [Chapter 2](#).

When you run the Installer to record a response file, it checks the system to verify that it meets the requirements to install the software. For this reason, Oracle recommends that you complete all of the required pre-installation tasks and record the response file while completing an installation.

2. If you have not installed Oracle software on this system previously, create the `oraInst.loc` file, as described in the previous section.
3. Make sure that the Oracle software owner user (typically `oracle`) has permissions to create or write to the Oracle home path that you will specify when you run the Installer.
4. To record a response file, enter a command similar to the following to start the Installer:

Note: Do not specify a relative path to the response file. If you specify a relative path, the Oracle Universal Installer fails.

```
$ /directory_path/runInstaller -record -destinationFile filename
```

In the previous example:

- *directory_path* is either the CD-ROM mount point directory, the path of the db directory on the DVD-ROM, or the path of the Disk1 directory on the hard drive
 - The `-record` parameter specifies that you want to record the responses that you enter in a response file
 - *filename* is the full path and file name of the response file that you want to record
5. On each Installer screen, specify the required information.
 6. When the Installer displays the Summary screen, do one of the following:
 - Click **Install** to create the response file, then continue with the installation.
 - Click **Cancel**, then **Yes** to create the response file but exit from the Installer without installing the software.

The response file is saved in the location that you specified using the `-destinationFile` option.

7. If you did not complete the installation, delete the Oracle home directory that the Installer created using the path you specified on the Specify File Locations screen.
8. Before using the recorded response file on another system, use any text editor to edit the file and make any required changes.

Use the comments in the file as a guide when editing it.

Run the Installer in Silent or Suppressed Mode

To run the Oracle Universal Installer in silent or suppressed mode, follow these steps:

1. Complete the pre-installation tasks listed in [Chapter 2](#).
2. Log in as the Oracle software owner user (typically `oracle`).
3. If you are completing a suppressed-mode installation, set the `DISPLAY` environment variable.

Note: You do not have to set the `DISPLAY` environment variable if you are completing a silent-mode installation.

4. To start the Installer in silent or suppressed mode, enter a command similar to the following:

Note: Do not specify a relative path to the response file. If you specify a relative path, the Oracle Universal Installer fails.

- Solaris:

```
$ /directory_path/runInstaller [-silent] -responseFile filename \  
  [-paramFile oraparamsilent.ini]
```

- Other operating systems:

```
$ /directory_path/runInstaller [-silent] -responseFile filename
```

In these examples:

- *directory_path* is either the CD-ROM mount point directory, the path of the db directory on the DVD-ROM, or the path of the Disk1 directory on the hard drive
- The `-silent` option indicates that you want to run the Installer in silent mode
- *filename* is the full path and filename of the installation response file that you configured
- The optional `-paramFile` option on Solaris specifies an alternative parameter file (`oraparamsilent.ini`) that you can use when performing silent installations. This alternative parameter file does not include checks for CDE or X Window System software patches for Solaris.

Note: For more information about other options for the `runInstaller` command, enter the following command:

```
$ /directory_path/runInstaller -help
```

5. When the installation completes, log in as the `root` user and run the `root.sh` script:

```
$ su
# /oracle_home_path/root.sh
```

Run NetCA in Silent Mode

You can run NetCA in silent mode to configure and start an Oracle Net listener on your system, configure naming methods, and configure Oracle Net service names. To run NetCA in silent mode, you must copy and edit a response file template. Oracle provides a response file template named `netca.rsp` in the `response` directory on the CD-ROM or in the `db/response` directory on the DVD-ROM.

Note: If you copied the software to a hard disk, the response file template is located in the `Disk1/response` directory.

To run NetCA in silent mode, follow these steps:

1. Copy the `netca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/response/netca.rsp local_directory
```

In this example, *directory_path* is the CD-ROM mount point directory or the db directory on the DVD-ROM. If you have copied the software to a hard drive, you can edit the file in the `response` directory if you prefer.

2. Open the response file in a text editor:

```
$ vi /local_dir/netca.rsp
```

3. Edit the file, following the instructions in the file.

Note: NetCA fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user (typically `oracle`) and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.
5. Enter a command similar to the following to run NetCA in silent mode:

```
$ $ORACLE_HOME/bin/netca /silent /responsefile=/local_dir/netca.rsp
```

In this example:

- The `/silent` option indicates that you want to run NetCA in silent mode
- `local_dir` is the full path of the directory where you copied the `netca.rsp` response file template

Run DBCA in Non-interactive or Silent Mode

You can run DBCA in non-interactive or silent mode to configure and start an Oracle Database on your system. To run DBCA in non-interactive or silent mode, you must copy and edit a response file template. Oracle provides a response file template named `dbca.rsp` in the `response` directory on the CD-ROM or in the `db/response` directory on the DVD-ROM.

Note: If you copied the software to a hard disk, the response file template is located in the `Disk1/response` directory.

Using DBCA in Non-interactive Mode

In non-interactive mode, DBCA uses values that you specify, in the response file or as command line options, to create a database. As it configures and starts the database, it displays a window that contains status messages and a progress bar. The window that it displays is the same window that is displayed when you choose to create a preconfigured database during an Enterprise Edition or Standard Edition installation.

To run DBCA in non-interactive mode, you must use a graphical display and set the `DISPLAY` environment variable.

Using DBCA in Silent Mode

In silent mode, DBCA uses values that you specify, in the response file or as command line options, to create a database. DBCA does not display any screens or information while it creates the database.

Running DBCA in Non-interactive or Silent Mode

To run DBCA in non-interactive or silent mode, follow these steps:

Note: As an alternative to editing the response file template, you can also create a database by specifying all required information as command line options when you run DBCA. For information about the list of options supported, enter the following command:

```
$ $ORACLE_HOME/bin/dbca -help
```

1. Copy the `dbca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/response/dbca.rsp local_directory
```

In this example, `directory_path` is the CD-ROM mount point directory or the `db` directory on the DVD-ROM. If you have copied the software to a hard drive, you can edit the file in the `response` directory if you prefer.

2. Open the response file in a text editor:

```
$ vi /local_dir/dbca.rsp
```

3. Edit the file, following the instructions in the file.

Note: DBCA fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user (typically `oracle`) and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.
5. If you intend running DBCA in non-interactive mode, set the `DISPLAY` environment variable.
6. Enter a command similar to the following to run DBCA in non-interactive or silent mode with a response file:

```
$ $ORACLE_HOME/bin/dbca {-progressOnly | -silent} [-cloneTemplate] \  
[-datafileDestination /datafilepath] -responseFile /local_dir/dbca.rsp
```

In this example:

- The `-silent` option indicates that you want to run DBCA in silent mode
- The `-progressOnly` option indicates that you want to run DBCA in non-interactive mode
- The `-cloneTemplate` option is required if you want to create one of the following preconfigured databases:
 - General Purpose
 - Transaction Processing
 - Data Warehouse
- `-datafileDestination` specifies the parent directory where you want to create the database files. If you do not specify this option, the parent directory for database files is the `$ORACLE_BASE/oradata` directory if the `ORACLE_BASE` environment variable is set, or the `$ORACLE_HOME/oradata` directory if `ORACLE_BASE` is not set.
- `local_dir` is the full path of the directory where you copied the `dbca.rsp` response file template

Using NAS Devices

If you have a network attached storage (NAS) device that has been certified through the Oracle Storage Compatibility Program (OSCP), you can use it to store the Oracle software, the Oracle database files, or both. This appendix provides guidelines for using a NAS storage device for Oracle software and database files. It includes information about the following:

- [General Configuration Guidelines for NAS Devices](#)
- [Choosing Mount Points](#)
- [Creating Files on a NAS Device for Use with ASM](#)
- [NFS Mount Options](#)

General Configuration Guidelines for NAS Devices

See the documentation provided with your NAS device for specific information about how to configure it. In addition, use the following guidelines to ensure that the performance of the Oracle software meets your requirements:

- Before using the NAS device for the installation, verify that it is certified. To identify certified devices, check the Oracle Storage Compatibility Program list at the following Web site:

<http://www.oracle.com/technology/deploy/availability/htdocs/oscp.html>

Note: The OSCP certifies NAS devices only for single-instance databases. For Oracle Real Application Clusters installations, see the Certify page on the Oracle*Metalink* Web site for information about certified storage solutions, including NAS:

<http://metalink.oracle.com>

- The performance of Oracle software and databases stored on NAS devices depends on the performance of the network connection between the Oracle server and the NAS device.

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

- For single-system installations (as opposed to RAC installations), you must create a separate Oracle home directory for each installation. Run the software in this Oracle home directory only from the system that you used to install it.

For Oracle Real Application Clusters (RAC) installations, you can use a single Oracle home directory mounted from each node in the cluster. You must mount this Oracle home directory on each node so that it has the same directory path on all nodes.

Choosing Mount Points

This section provides guidelines on how to choose the mount points for the file systems that you want to use for the Oracle software and database files. The guidelines contained in this section comply with the Optimal Flexible Architecture recommendations. It contains information about the following:

- [Choosing Mount Points for Oracle Software Files](#)
- [Choosing Mount Points for Oracle Database and Recovery Files](#)

Choosing Mount Points for Oracle Software Files

Oracle software files are stored in three different directories:

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

For the first installation of Oracle software on a system, the Oracle base directory, identified by the ORACLE_BASE environment variable, is normally the parent

directory for both the Oracle Inventory and Oracle home directories. For example, for a first installation, the Oracle base, Oracle Inventory, and Oracle home directories might have paths similar to the following:

Directory	Path
Oracle base (\$ORACLE_BASE)	/u01/app/oracle
Oracle Inventory	\$ORACLE_BASE/oraInventory
Oracle home	\$ORACLE_BASE/product/10.1.0/db_1

For subsequent installations, you can choose to use either the same Oracle base directory or a different one, but every subsequent installation uses the original Oracle Inventory directory. For example, if you use the /u02/app/oracle directory as the Oracle base directory for a new installation, the Oracle Inventory directory continues to be /u01/app/oracle/oraInventory.

To enable you to effectively maintain the Oracle software on a particular system, Oracle recommends that you locate the Oracle Inventory directory only on a local file system, if possible. If you must place the Oracle Inventory directory on a NAS device, create a specific directory for each system, to prevent more than one system from writing to the same Inventory.

Directory-Specific Guidelines

You can use any of the following directories as mount points for NFS file systems used to store Oracle software:

Note: In the following examples, the paths shown are the defaults if the ORACLE_BASE environment variable is set before you start the Installer.

- Oracle base directory or its parents (/u01/app/oracle for example)

If you use the Oracle base directory or one of its parents as a mount point, the default location for all Oracle software and database files will be on that file system. During the installation, you might consider changing the default location of the following directories:

- The Oracle Inventory directory (*oracle_base/oraInventory*)

Specify a local file system or a host-specific directory on the NFS file system, for example:

```
oracle_base/hostname/oraInventory
```

- The Oracle database file directory (*oracle_base/oradata*)

You might want to use a different file system for database files, for example, to enable you to specify different mount options or to distribute I/O.

- The Oracle database recovery file directory (*oracle_base/flash_recovery_area*)

Oracle recommends that you use different file systems for database and recovery files.

If you use this mount point, all Oracle installations that use this Oracle base directory will use the NFS file system.

- The product directory (*oracle_base/product*)

By default, only software files will be located on the NFS file system. You can also use this mount point to install software from different releases, for example:

```
/u01/app/oracle/product/9.2.0
/u01/app/oracle/product/10.1.0/db_1
```
- The release directory (*oracle_base/product/10.1.0*)

By default, only software files will be located on the NFS file system. You can also use this mount point to install different products from the same release, for example:

```
/u01/app/oracle/product/10.1.0/crs
/u01/app/oracle/product/10.1.0/db_1
/u01/app/oracle/product/10.1.0/companion_1
```
- The Oracle home directory (*oracle_base/product/10.1.0/db_1*)

By default, only software files will be located on the NFS file system. This is the most restrictive mount point. You can use it only to install a single release of one product:

```
/u01/app/oracle/product/10.1.0/db_1
```

Choosing Mount Points for Oracle Database and Recovery Files

To store Oracle database or recovery files on a NAS device, you can use different paths depending on whether you want to store files from only one database or from more than one database:

- Use the NFS file system for files from more than one database

If you want to store the database files or recovery files from more than one database on the same NFS file systems, use paths or mount points similar to the following:

File Type	Path or Mount Point
Database files	/u02/oradata
Recovery files	/u03/flash_recovery_area

When the Installer prompts you for the datafile and the recovery file directories, specify these paths. The DBCA and Enterprise Manager create subdirectories in these directories using the value you specify for the database name (DB_NAME) as the directory name, for example:

```
/u02/oradata/db_name1
/u03/flash_recovery_area/db_name1
```

- Use the NFS file system for files from only one database

If you want to store the database files or recovery files for only one database in the NFS file system, you can create mount points similar to the following, where *orcl* is the name that you want to use for the database:

```
/u02/oradata/orcl
/u03/flash_recovery_area/orcl
```

Specify the directory `/u02/oradata` when the Installer prompts you for the datafile directory and specify the directory `/u03/flash_recovery_area` when the Installer prompts you for the recovery file location. The `orcl` directory will be used automatically either by DBCA or by Enterprise Manager.

Creating Files on a NAS Device for Use with ASM

If you have a certified NAS storage device, you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an ASM disk group. To create these files, follow these steps:

Note: To use files as disk devices in an ASM disk group, the files must be on an NFS mounted file system. You cannot use files on local file systems.

1. If necessary, create an exported directory for the disk group files on the NAS device.

See the NAS device documentation for more information about completing this step.

2. Switch user to `root`:

```
$ su -
```

3. Create a mount point directory on the local system:

```
# mkdir -p /mnt/oracleasm
```

4. To ensure that the NFS file system is mounted when the system reboots, add an entry for the file system to the appropriate mount file for your operating system (`/etc/filesystems` on AIX, `/etc/vfstab` on Solaris, and `/etc/fstab` on other platforms).

For more information about editing the mount file for your operating system, see the man pages. For more information about recommended mount options, see the ["NFS Mount Options"](#) section on page C-6.

5. Enter a command similar to the following to mount the NFS file system on the local system:

```
# mount /mnt/oracleasm
```

6. Choose a name for the disk group that you want to create, for example `nfsdg`.

7. Create a directory for the files on the NFS file system, using the disk group name as the directory name:

```
# mkdir /mnt/oracleasm/nfsdg
```

8. Use commands similar to the following, depending on your operating system, to create the required number of zero-padded files in this directory:

- Solaris:

```
# mkfile 1024M /mnt/oracleasm/nfsdg/disk1
```

- Other operating systems:

```
# dd if=/dev/zero of=/mnt/oracleasm/nfsdg/disk1 bs=1024k count=1000
```

Both examples create 1 GB files on the NFS file system. You must create one, two, or three files respectively to create an external, normal, or high redundancy disk group.

- Enter the following commands to change the owner, group, and permissions on the directory and files that you created:

```
# chown -R oracle:dba /mnt/oracleasm
# chmod -R 660 /mnt/oracleasm
```

- When you are creating the database, edit the ASM disk discovery string to specify a regular expression that matches the file names you created. For example, you might specify a disk discovery string similar to the following:

```
/mnt/oracleasm/nfsdg/*
```

NFS Mount Options

When mounting an NFS file system on your system, Oracle recommends that you use the same mount point options that your NAS vendor used when certifying the device. See your device documentation or contact your vendor for information about recommended mount-point options.

In general, most vendors recommend that you use the NFS mount options listed in the following table. See your operating system or NAS device documentation for more information about the specific options supported on your platform.

Option	Description
hard	Generate a hard mount of the NFS file system. If the connection to the server is lost temporarily, Oracle continues to retry the connection until the NAS device responds.
bg	Try to connect in the background if connection fails.
proto=tcp (or tcp on Linux)	Use the TCP protocol rather than UDP. TCP is more reliable than UDP.
vers=3 (or nfsvers=3 on Linux)	Use NFS version 3. Oracle recommends that you use NFS version 3 where available, unless the performance of version 2 is higher.
suid	Allow clients to run executables with SUID enabled. This option is required for Oracle software mount points.
rsize, wsize	The number of bytes used when reading or writing to the NAS device. A value of 8192 is often recommended for NFS version 2 and 32768 is often recommended for NFS version 3.
nointr (or intr)	Do not allow (or allow) keyboard interrupts to kill a process that is hung while waiting for a response on a hard-mounted file system. Note: Different vendors have different recommendations about this option. Contact your vendor for advice.
noac	Disable attribute caching. Note: You must specify this option for NFS file systems where you want to install the software. If you do not use this option, the Installer will not install the software in the directory that you specify.
forcedirectio (Solaris)	Use direct I/O (no buffering) for data transfers. Note: Use this option for file systems that contain only Oracle database files. Do not use it for a file system that contains Oracle software.

Optimal Flexible Architecture

This appendix describes the Optimal Flexible Architecture (OFA) standard. The OFA standard is a set of configuration guidelines created to ensure reliable Oracle installations that require little maintenance. It includes information about the following topics:

- [Changes to the Optimal Flexible Architecture for Oracle Database 10g](#)
- [Overview of the Optimal Flexible Architecture Standard](#)
- [Optimal Flexible Architecture Implemented on UNIX Systems](#)

Changes to the Optimal Flexible Architecture for Oracle Database 10g

For previous releases of Oracle Database, the OFA recommended Oracle home path was similar to the following:

```
/u01/app/oracle/product/9.2.0
```

For Oracle Database 10g, the OFA recommended Oracle home path has changed. The OFA recommended path is now similar to the following:

```
/u01/app/oracle/product/10.1.0/type[_n]
```

In this example, *type* is the type of Oracle home, for example Oracle Database (*db*) or Oracle Client (*client*), and *n* is an optional counter. This syntax provides the following benefits:

- You can install different products with the same release number in the same Oracle base directory, for example:

```
/u01/app/oracle/product/10.1.0/db_1  
/u01/app/oracle/product/10.1.0/client_1
```

- You can install the same product more than once in the same Oracle base directory, for example:

```
/u01/app/oracle/product/10.1.0/db_1  
/u01/app/oracle/product/10.1.0/db_2
```

Overview of the Optimal Flexible Architecture Standard

The OFA standard is designed to:

- Organize large amounts of complicated software and data on disk, to avoid device bottlenecks and poor performance
- Facilitate routine administrative tasks such as software and data backup, which are often vulnerable to data corruption
- Facilitate switching between multiple Oracle databases
- Adequately manage and administer database growth
- Help eliminate fragmentation of free space in the data dictionary, isolate other fragmentation, and minimize resource contention

Oracle recommends that you implement the OFA standard when you install and configure Oracle products.

Characteristics of an OFA-Compliant Installation

This section describes the characteristics of an Oracle product installation that complies with the OFA standard.

File System Organization

The file system is organized to enable easy administration and to facilitate:

- Adding data into existing databases
- Adding users
- Creating databases
- Adding hardware

Distributed I/O Loads

I/O loads are distributed across enough disk drives to prevent performance bottlenecks.

Hardware Support

In most cases, investment in new hardware is not required to take advantage of the OFA standard.

Safeguards Against Drive Failures

By distributing applications across more than one drive, drive failures affect as few applications as possible.

Distribution of Home Directories

The following items can be distributed across more than one disk drive:

- The collection of home directories
- The contents of an individual home directory

Integrity of Login Home Directories

You can add, move, or delete login home directories without having to revise programs that refer to them.

Independence of UNIX Directory Subtrees

Categories of files are separated into independent UNIX directory subtrees so that files in one category are minimally affected by operations on files in other categories.

Supports Concurrent Execution of Application Software

You can run multiple versions of Oracle software simultaneously, enabling you to test and use a new release before retiring the previous release. Transferring to a new release after an upgrade is simple for the administrator and transparent for the user.

Separates Administrative Information for Each Database

The ability to separate administrative information for each database ensures a reasonable structure for the organization and storage of administrative data.

Uses Consistent Database File Naming

Database files are named so that:

- Database files are easy to distinguish from other files
- Files belonging to one database are easy to distinguish from files that belong to another database
- Control files, redo log files, and datafiles can be identified as such
- The association of datafile to tablespace is clearly indicated

Separation of Tablespace Contents

Tablespace contents are separated to:

- Minimize tablespace free space fragmentation
- Minimize I/O request contention
- Maximize administrative flexibility

I/O Loads Tuned Across All Drives

I/O loads are tuned across all drives, including drives storing Oracle data in either Automatic Storage Management disk groups or in raw devices.

Optimal Flexible Architecture Implemented on UNIX Systems

This section describes the naming strategy recommended by the OFA standard for UNIX systems.

File Systems

This section describes the naming conventions for mount points.

Number of File Systems

To fully implement the OFA recommendations for a database stored on file systems that are not striped or mirrored, you require at least three file systems located on separate physical devices.

Naming Conventions

Name all file system mount points using the syntax */pm*, where *p* is a string constant and *m* is a unique fixed-length key (typically a two-digit number) used to distinguish each mount point. For example: */u01* and */u02*, or */disk01* and */disk02*.

Naming Mount Points for Very Large Databases (VLDBs)

If each disk drive contains database files from one application and there are enough drives for each database to prevent I/O bottlenecks, use the syntax */pm/q/dm* for naming mount points. [Table D-1](#) describes the variables used in this syntax.

Table D–1 Syntax for Naming Mount Points for Very Large Databases

Variable	Description
<i>pm</i>	A mount point name
<i>q</i>	A string denoting that Oracle data is stored in this directory, for example, <code>oradata</code>
<i>dm</i>	The value of the initialization parameter <code>DB_NAME</code> (typically the same as the instance SID for single-instance databases)

For example, to allocate two drives exclusively for the `test` database, name the mount points `/u01/oradata/test` and `/u02/oradata/test`.

Naming Directories

This section describes the naming conventions for OFA compliant directories.

Oracle Base Directory Naming Convention

The Oracle base directory is the top-level directory for Oracle products installed by the same user. Name Oracle base directories using the syntax `/pm/h/u`. [Table D–2](#) describes the variables used in this syntax.

Table D–2 Syntax for Naming Oracle Base Directories

Variable	Description
<i>pm</i>	A mount point name
<i>h</i>	A standard directory name
<i>u</i>	The name of the owner of the directory (the user running the Installer)

For example, `/u01/app/oracle` is an Oracle base directory created by the `oracle` user and `/u01/app/appmgr` is an Oracle base directory created by the `appmgr` user.

Placing Oracle base directories at the same level in the UNIX file system is advantageous because it allows you to refer to the collection of Oracle base directories on different mount points using a single pattern matching string, `/*/app/*`.

Referring to Pathnames

Refer to explicit pathnames only in files designed specifically to store them, such as the password file, `/etc/passwd`, and the Oracle `oratab` file. Refer to group memberships only in the `/etc/group` file.

Oracle Home Directory Naming Convention

To help fulfill the OFA requirement of simultaneously running multiple versions of Oracle software, install the software in a directory matching the pattern `/pm/h/u/product/v/type_[n]`.

[Table D–3](#) describes the variables used in this syntax.

Table D–3 Syntax for Naming Oracle Home Directories

Variable	Description
<i>pm</i>	A mount point name
<i>h</i>	A standard directory name
<i>u</i>	The name of the owner of the directory
<i>v</i>	The version of the software
<i>type</i>	The type of installation, for example Database (<i>db</i>), Client (<i>client</i>), Companion (<i>companion</i>), or CRS (<i>crs</i>)
<i>n</i>	An optional counter, which enables you to install the same product more than once in the same Oracle base directory

For example:

- `/u01/app/oracle/product/10.1.0/db_1` indicates the Oracle home directory for the first installation of Oracle Database on this system
- `/u01/app/oracle/product/10.1.0/crs` indicates the Oracle home directory for Oracle Cluster Ready Services (CRS is required for RAC installations)
Oracle CRS can be installed only once on the system, so the optional counter is not required.

Set the `ORACLE_HOME` environment variable after installation to specify the Oracle home directory.

Naming Subdirectories

To facilitate the organization of administrative data, Oracle recommends that you store database-specific administration files in subdirectories matching the pattern `/h/admin/d/a/`, where *h* is the Oracle base directory, *d* is the database name (DB_NAME), and *a* is a subdirectory for specific types of database administration files. [Table D–4](#) describes the database administration file subdirectories.

Table D–4 Subdirectories for Database Administration Files

Subdirectory	Description
<code>adhoc</code>	Ad hoc SQL scripts
<code>arch</code>	Archived redo log files
<code>adump</code>	Audit files (Set the <code>AUDIT_FILE_DEST</code> initialization parameter to specify the <code>adump</code> directory. Clean out this subdirectory periodically.)
<code>bdump</code>	Background process trace files
<code>cdump</code>	Core dump files
<code>create</code>	Scripts used to create the database
<code>exp</code>	Database export files
<code>logbook</code>	Files recording the status and history of the database
<code>pfile</code>	Instance parameter files
<code>udump</code>	User SQL trace files

For example, `/u01/app/oracle/admin/sab/adhoc/` is the `adhoc` subdirectory associated with the database named `sab`.

Naming Database Files

The following table lists the recommended file naming conventions for database files:

Note: Oracle Managed Files (OMF) and files stored in ASM disk groups use different naming conventions. For more information about these naming conventions, refer to the *Oracle Database Administrator's Guide*.

File Type	File Naming Convention
Control files	<code>/pm/q/d/control.ctl</code>
Redo log files	<code>/pm/q/d/redon.log</code>
Datafiles	<code>/pm/q/d/tn.dbf</code>

The following table describes this syntax:

Variable	Description
<i>pm</i>	A mount point name described previously in this appendix
<i>q</i>	A string (typically <code>oradata</code>) distinguishing Oracle data from all other files
<i>d</i>	The value of the <code>DB_NAME</code> initialization parameter (typically, the same as the instance <code>SID</code> for single-instance databases)
<i>t</i>	An Oracle tablespace name
<i>n</i>	A two-digit string

Note: Do not store files other than control files, redo log files, or datafiles associated with database *d* in the path `/pm/q/d`.

Using this convention, it is easy to determine the database to which the `/u03/oradata/sab/system01.dbf` file belongs.

Separating Segments With Different Requirements

Separate groups of segments with different lifespans, I/O request demands, and backup frequencies across different tablespaces.

[Table D-5](#) describes the special tablespaces that the Database Configuration Assistant creates for each Oracle database. If you manually create a database, you must create the required tablespaces. These tablespaces are in addition to those required for application segments.

See Also: See the *Oracle Database Administrator's Guide* for information about creating databases manually.

Table D-5 Special Tablespaces

Tablespace	Required	Description
EXAMPLE	No	The EXAMPLE tablespace used to store the Sample Schemas
SYSAUX	Yes	Auxiliary tablespace to the SYSTEM tablespace
SYSTEM	Yes	Data dictionary segments
TEMP	Yes	Temporary segments
UNDOTBS1	Yes	Used by Oracle to store undo information
USERS	No	Miscellaneous user segments

Creating these special tablespaces is effective because data dictionary segments are never dropped, and no other segments that can be dropped are allowed in the SYSTEM tablespace. Doing this ensures that the SYSTEM tablespace does not require a rebuild due to tablespace free-space fragmentation.

Naming Tablespaces

Name tablespaces descriptively using a maximum of eight characters. Although Oracle Database tablespace names can be 30 characters long, portable UNIX filenames are restricted to 14 characters. The recommended standard for a datafile basename is *tn.dbf*, where *t* is a descriptive tablespace name and *n* is a two-digit string. Because the extension and the two-digit string take six characters, only eight characters remain for the tablespace name.

Descriptive names enable the datafile to be associated with the tablespace that uses it. For example, the names GLD and GLX might be used for the tablespaces storing General Ledger data and General Ledger indexes, respectively.

Note: Do not embed reminders of the word "tablespace" in your tablespace names. Tablespace names can be distinguished by context. For example, do not name the General Ledger tablespace GLD_TBS01.dbf.

Exploiting the OFA Structure for Oracle Files

Table D-6 describes the syntax used for identifying classes of files.

Table D-6 Directory Structure Syntax for Identifying Classes of Files

Directory Structure Syntax	Description
/u[0-9][0-9]	User data directories
/home/	User home directories
/app/	User application software directories
*/app/applmgr	Oracle applications software subtrees
*/app/oracle/product	Oracle software subtrees
*/app/oracle/product/10.1.0	Oracle software subtree for release 10.1.0 products
/app/oracle/product/10.1.0/db	Oracle home directories for Oracle Database 10g
*/app/oracle/admin/sab	sab database administrative subtrees
/app/oracle/admin/sab/arch/	sab database archived log files

Table D–6 (Cont.) Directory Structure Syntax for Identifying Classes of Files

Directory Structure Syntax	Description
<code>*/oradata</code>	Oracle data directories
<code>*/oradata/sab/*</code>	sab database files
<code>*/oradata/sab/*.log</code>	sab database redo log files

OFA File Mapping

Table D–7 shows a hierarchical file mapping of a sample OFA-compliant installation with two Oracle home directories and two databases. The database files are distributed across three mount points, /u02, /u03, and /u04.

Table D–7 Hierarchical File Mapping for OFA Installation

Directory	Description
<code>/</code>	Root directory
<code>/u01/</code>	User data mount point 1
<code>/u01/app/</code>	Subtree for application software
<code>/u01/app/oracle/</code>	Oracle Base directory
<code>/u01/app/oracle/admin/</code>	Subtree for database administration files
<code>/u01/app/oracle/admin/TAR</code>	Subtree for support log files
<code>/u01/app/oracle/admin/db_name1/</code>	admin subtree for <i>db_name1</i> database
<code>/u01/app/oracle/admin/db_name2/</code>	admin subtree for <i>db_name2</i> database
<code>/u01/app/oracle/doc/</code>	Online documentation
<code>/u01/app/oracle/flash_recovery_area/</code>	Subtree for recovery files
<code>/u01/app/oracle/flash_recovery_area/db_name1</code>	Recovery files for <i>db_name1</i> database
<code>/u01/app/oracle/flash_recovery_area/db_name2</code>	Recovery files for <i>db_name2</i> database
<code>/u01/app/oracle/product/</code>	Distribution files
<code>/u01/app/oracle/product/9.2.0</code>	Oracle home directory for Oracle9i release 2 (9.2.0)
<code>/u01/app/oracle/product/10.1.0/db_1</code>	Oracle home directory for Oracle Database 10g release 1 (10.1.0)
<code>/u01/app/kjff/</code>	Oracle base directory for user <i>kjff</i>
<code>/u01/app/edm/</code>	Oracle base directory for user <i>edm</i>
<code>/u02</code>	User data mount point 2
<code>/u02/oradata/</code>	Subtree for Oracle data
<code>/u02/oradata/db_name1/</code>	Subtree for <i>db_name1</i> database files
<code>/u02/oradata/db_name2/</code>	Subtree for <i>db_name2</i> database files
<code>/u03/</code>	User data mount point 3
<code>/u03/oradata/</code>	Subtree for Oracle data
<code>/u03/oradata/db_name1/</code>	Subtree for <i>db_name1</i> database files
<code>/u03/oradata/db_name2/</code>	Subtree for <i>db_name2</i> database files
<code>/u04/</code>	User data mount point 4

Table D–7 (Cont.) Hierarchical File Mapping for OFA Installation

Directory	Description
/u04/oradata/	Subtree for Oracle data
/u04/oradata/db_name1/	Subtree for <i>db_name1</i> database files
/u04/oradata/db_name2/	Subtree for <i>db_name2</i> database files

Default Port Numbers

During installation, the Installer assigns port numbers to components from a set of default port numbers. This appendix lists the default port numbers and describes how to change the assigned port after installation. It includes information about the following topics:

- [Components and Port Ranges](#)
- [Changing the Enterprise Manager Agent HTTP Port](#)
- [Changing the Enterprise Manager Database Control Ports](#)
- [Changing the iSQL*Plus Ports](#)
- [Changing the Ultra Search Ports](#)

Components and Port Ranges

The following table lists the port ranges used by components that are configured during the installation. By default, the first port in the range is assigned to the component, if it is available.

Component	Port	Port Range
Enterprise Manager Agent	HTTP	1830 - 1849
Enterprise Manager Database Control	HTTP	5500 - 5519
	RMI	5520 - 5539
	JMS	5540 - 5559
iSQL*Plus	HTTP	5560 - 5579
	RMI	5580 - 5599
	JMS	5600 - 5619
Ultra Search	HTTP	5620 - 5639
	RMI	5640 - 5659
	JMS	5660 - 5679

Changing the Enterprise Manager Agent HTTP Port

To change the Oracle Enterprise Manager Agent HTTP port, edit the following files:

- `$ORACLE_HOME/host_sid/sysman/config/emoms.properties`

Modify the following parameter in the file:

```
oracle.sysman.eml.mntr.repAgentUrl=http://host.domain:1830/emd/main
```

- `$ORACLE_HOME/host_sid/sysman/config/emd.properties`

Modify the following parameter in the file:

```
EMD_URL=http://host.domain:1830/emd/main
```

Changing the Enterprise Manager Database Control Ports

The following sections describe how to change the Oracle Enterprise Manager Database Control ports.

Changing the HTTP Port

To change the HTTP port, edit the following files:

- `$ORACLE_HOME/host_sid/sysman/config/emoms.properties`

Modify the following parameters in the file:

```
oracle.sysman.emSDK.svlt.ConsoleServerPort=5500
oracle.sysman.emSDK.svlt.ConsoleServerHTTPSPort=5500
```

- `$ORACLE_HOME/host_sid/sysman/config/emd.properties`

Modify the following parameters in the file:

```
REPOSITORY_URL=http://host.domain:5500/em/upload/
emdWalletSrcUrl=http://host.domain:5500/em/wallets/emd
```

- `$ORACLE_HOME/oc4j/j2ee/OC4J_DBCConsole_host_sid/config/http-web-site.xml`

Modify the port attribute of the web-site element:

```
<web-site port="5500" ...>
```

Changing the RMI Port

To change the RMI port, modify the port attribute of the rmi-server element in the `$ORACLE_HOME/oc4j/j2ee/OC4J_DBCConsole_host_sid/config/rmi.xml` file:

```
<rmi-server port="5520" ...>
```

Changing the JMS Port

To change the JMS port, modify the port attribute of the jms-server element in the `$ORACLE_HOME/oc4j/j2ee/OC4J_DBCConsole_host_sid/config/jms.xml` file:

```
<jms-server port="5540" ...>
```

Changing the iSQL*Plus Ports

The following sections describe how to change the iSQL*Plus ports.

Changing the HTTP Port

To change the HTTP port, edit the following files:

- `$ORACLE_HOME/host_sid/sysman/config/emoms.properties`

Modify the following parameters in the file:

```
oracle.sysman.db.isqlplusUrl=http://host.domain\:5560/isqlplus/dynamic
oracle.sysman.db.isqlplusWebDBAUrl=http://host.domain\:5560/isqlplus/dynamic
```

- `$ORACLE_HOME/oc4j/j2ee/isqlplus/config/http-web-site.xml`

Modify the port attribute of the web-site element:

```
<web-site port="5560" ...>
```

Changing the RMI Port

To change the RMI port, modify the port attribute of the rmi-server element in the `$ORACLE_HOME/oc4j/j2ee/isqlplus/config/rmi.xml` file:

```
<rmi-server port="5580" ...>
```

Changing the JMS Port

To change the JMS port, modify the port attribute of the jms-server element in the `$ORACLE_HOME/oc4j/j2ee/isqlplus/config/jms.xml` file:

```
<jms-server port="5600" ...>
```

Changing the Ultra Search Ports

The following sections describe how to change the Ultra Search ports.

Changing the HTTP Port

To change the HTTP port, modify the port attribute of the web-site element in the `$ORACLE_HOME/oc4j/j2ee/OC4J_SEARCH/config/http-web-site.xml` file:

```
<web-site port="5620"...>
```

Changing the RMI Port

To change the RMI port, modify the port attribute of the rmi-server element in the `$ORACLE_HOME/oc4j/j2ee/OC4J_SEARCH/config/rmi.xml` file:

```
<rmi-server port="5640"...>
```

Changing the JMS Port

To change the JMS port, modify the port attribute of the jms-server element in the `$ORACLE_HOME/oc4j/j2ee/OC4J_SEARCH/config/jms.xml` file:

```
<jms-server port="5660"...>
```

Troubleshooting

This appendix contains information about troubleshooting. It includes information about the following topics:

- [Verify Requirements](#)
- [X Windows Display Errors](#)
- [What to Do If an Installation Error Occurs](#)
- [Reviewing the Log of an Installation Session](#)
- [Troubleshooting Configuration Assistants](#)
- [Silent-Mode Response File Error Handling](#)
- [Cleaning Up After a Failed Installation](#)

Verify Requirements

Before performing any of the troubleshooting steps in this appendix, ensure that the system meets the requirements and that you have completed all of the pre-installation tasks specified in [Chapter 2, "Pre-installation Tasks"](#).

Read the Release Notes

Read the release notes for the product on your platform before installing it. The release notes are available on the Oracle Database 10g disc. The latest version of the release notes is also available on the OTN Web site:

<http://www.oracle.com/technology/documentation/>

X Windows Display Errors

If you are running the Installer on a remote system and you want to display the Installer's user interface on your local system, you might see error messages similar to the following:

```
"Failed to connect to server"  
"Connection refused by server"  
"Can't open display"
```

If you see one of these error messages, follow these steps:

Note: This procedure applies only to users of UNIX workstations. If you are using a PC or other system with X server software installed, see the X server documentation for information about how to permit remote systems to display X applications on the local system.

1. In a local terminal window, log in as the user that started the X window session.
2. Enter the following command:

```
$ xhost +
```
3. Enter the following commands, where *workstation_name* is the host name or IP address of your workstation:
 - Bourne, Bash, or Korn shell:

```
$ DISPLAY=workstation_name:0.0  
$ export DISPLAY
```
 - C shell:

```
% setenv DISPLAY workstation_name:0.0
```
4. To determine whether X Window applications display correctly on the local system, enter the following command:

```
$ xclock
```

The X clock should appear on your monitor.
5. If the X clock appears, close X clock and start the Installer again.

What to Do If an Installation Error Occurs

If you encounter an error during installation:

- Do not exit the Installer.
- If you clicked **Next** after you entered incorrect information on one of the installation screens, click **Back** to return to the screen and correct the information.
- If you encounter an error while the Installer is copying or linking files, see the ["Reviewing the Log of an Installation Session"](#) section on page F-3.
- If you encounter an error while a configuration assistant is running, see the ["Troubleshooting Configuration Assistants"](#) section on page F-4.
- If you cannot resolve the problem, remove the failed installation by following the steps listed in the ["Cleaning Up After a Failed Installation"](#) section on page F-5.

Reviewing the Log of an Installation Session

During an installation, the Installer records all of the actions that it performs in a log file. If you encounter problems during the installation, review the log file for information about possible causes of the problem.

To view the log file, follow these steps:

1. If necessary, enter one of the following commands to determine the location of the `oraInventory` directory:

- AIX, Linux x86, or Linux Itanium:

```
$ cat /etc/oraInst.loc
```

- Other operating systems:

```
$ cat /var/opt/oracle/oraInst.loc
```

The `inventory_loc` parameter in this file specifies the location of the `oraInventory` directory.

2. Enter the following command to change directory to the Installer log file directory, where `orainventory_location` is the location of the `oraInventory` directory:

```
$ cd /orainventory_location/logs
```

3. Enter the following command to determine the file name of the log file:

```
$ ls -ltr
```

This command lists the files in the order of creation, with the most recent file shown last. Installer log files have names similar to the following, where `date_time` indicates the date and time that the installation started:

```
installActionsdate_time.log
```

4. To view the most recent entries in the log file, where information about a problem is most likely to appear, enter a command similar to the following:

```
$ tail -50 installActionsdate_time.log | more
```

This command displays the last 50 lines in the log file.

5. If the error displayed by the Installer or listed in the log file indicates a relinking problem, see the following file for more information:

`$ORACLE_HOME/install/make.log`

Troubleshooting Configuration Assistants

To troubleshoot an installation error that occurs when a configuration assistant is running:

- Review the installation log files listed in the ["Reviewing the Log of an Installation Session"](#) section on page F-3.
- Review the specific configuration assistant log file located in the `$ORACLE_HOME/cfgtoollogs` directory. Try to fix the issue that caused the error.
- If you see the "Fatal Error. Reinstall" message, look for the cause of the problem by reviewing the log files. Refer to the ["Fatal Errors"](#) section for further instructions.

Configuration Assistant Failure

Oracle configuration assistant failures are noted at the bottom of the installation screen. The configuration assistant interface displays additional information, if available. The configuration assistant execution status is stored in the following file:

`oraInventory_location/logs/installActionsdate_time.log`

The execution status codes are listed in the following table:

Status	Result Code
Configuration assistant succeeded	0
Configuration assistant failed	1
Configuration assistant cancelled	-1

Fatal Errors

If you receive a fatal error while a configuration assistant is running, you must remove the current installation and reinstall the Oracle software, as follows:

1. Remove the failed installation as described in the ["Cleaning Up After a Failed Installation"](#) section on page F-5.
2. Correct the cause of the fatal error.
3. Reinstall the Oracle software.

Silent-Mode Response File Error Handling

To determine whether a silent-mode installation succeeds or fails, see the following log file:

`/oraInventory_location/logs/silentInstalldate_time.log`

If necessary, see the previous section for information about determining the location of the `oraInventory` directory.

A silent installation fails if:

- You do not specify a response file
- You specify an incorrect or incomplete response file
- The Installer encounters an error, such as insufficient disk space

The Installer or configuration assistant validates the response file at runtime. If the validation fails, the silent-mode installation or configuration process ends. The Installer treats values for parameters that are of the wrong context, format, or type as if no value was specified in the file.

Cleaning Up After a Failed Installation

If an installation fails, you must remove files that the Installer created during the attempted installation and remove the Oracle home directory. Perform the following steps to remove the files:

1. Start the Installer as described in the ["Install the Oracle Database Software"](#) section on page 3-5.
2. Click **Deinstall Products** on the Welcome window or click **Installed Products** on any Installer window.

The Inventory window appears, listing installed products.

3. Select the Oracle home that contains the products that you want to remove, then click **Remove**.
4. Manually remove the Oracle home directory created during the failed installation.
5. Reinstall the Oracle software.

Numerics

- 64-bit
 - checking system architecture, 2-5

A

- accounts
 - reviewing, 5-3
 - unauthenticated access to, 5-7
- Adaptive Server
 - See Sybase Adaptive Server Enterprise
- Advanced Security Option
 - See Oracle Advanced Security
- AdvFS
 - identifying AdvFS file systems on Tru64 UNIX, 2-65
- aio_task_max_num subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- AIX
 - APAR download location, 2-8
 - character raw device names, 2-53
 - checking disk availability for ASM, 2-53, 2-57, 2-60, 2-77, 2-79
 - checking disk availability for raw devices, 2-67
 - checking existence of nobody user, 2-31
 - checking maintenance level, 2-7
 - checking patch requirements, 2-8
 - checking software requirements, 2-7
 - checking version, 2-7
 - configuring disks for ASM, 2-53
 - configuring new disks, 2-53, 2-67
 - configuring raw devices, 2-67
 - creating a raw device mapping file, 2-70
 - creating a volume group, 2-67
 - creating logical volumes, 2-69
 - creating raw logical volumes, 2-68
 - creating volume groups, 2-68
 - identifying disks for ASM, 2-53
 - identifying disks for LVM, 2-67
 - identifying LVM disks, 2-53, 2-67
 - initializing disks for LVM, 2-67
 - location of oratab file, 2-44
 - patch requirements, 2-7
 - PTF download location, 2-8
 - raw device sizes, 2-68
 - software requirements, 2-6
 - specifying owner and permissions of raw devices, 2-70
- AIXwindows
 - Motif library requirement, 2-6
- AL24UTFFSS character sets, 1-11
- ANONYMOUS user
 - unauthenticated account access with, 5-7
- APAR
 - checking, 2-8
 - download location, 2-8
 - requirements on AIX, 2-7
- architecture
 - checking system architecture, 2-5
- ASM
 - block and character device names
 - on HP-UX, 2-54
 - on Solaris, 2-64, 2-85
 - on Tru64 UNIX, 2-65
 - block device names on Linux, 2-57, 2-60
 - changing owner and permissions of disks
 - on AIX, 2-53
 - on HP-UX, 2-54
 - on Solaris, 2-64
 - on Tru64 UNIX, 2-66
 - changing owner and permissions of raw devices
 - on Linux, 2-62, 2-63, 2-82
 - character raw device names on AIX, 2-53
 - characteristics of failure groups, 2-51
 - checking disk availability on AIX, 2-53, 2-57, 2-60, 2-77, 2-79
 - checking disk availability on HP-UX, 2-54
 - checking disk availability on Solaris, 2-64, 2-85
 - checking disk availability on Tru64 UNIX, 2-65
 - configuring disks
 - on AIX, 2-53
 - on HP-UX, 2-54
 - on Linux, 2-55
 - on Solaris, 2-63
 - on Tru64 UNIX, 2-65
 - disk groups, 2-50
 - displaying attached disks
 - on Linux, 2-57, 2-60, 2-77, 2-79
 - failure groups, 2-50
 - examples, 2-51
 - identifying, 2-51

- identifying available disks on AIX, 2-53
- identifying available disks on HP-UX, 2-54
- identifying available disks on Linux, 2-57, 2-60, 2-77, 2-79
- identifying available disks on Solaris, 2-64, 2-85
- identifying available disks on Tru64 UNIX, 2-65
- identifying disks on AIX, 2-53
- identifying disks on HP-UX, 2-54
- identifying disks on Linux, 2-57, 2-60, 2-77, 2-79
- identifying disks on Solaris, 2-64, 2-85
- identifying disks on Tru64 UNIX, 2-65
- OFA file naming conventions, D-7
- overview, 1-6
- recommendations for disk groups, 2-50
- redundancy levels, 1-6
- space required for preconfigured database, 2-50
- storage option for datafiles, 2-46
- using NAS files, C-5

authorized problem analysis report
See APAR

B

- base directory
 - See* Oracle base directory
- Bash shell
 - default user startup file, 2-93
 - setting shell limits on Linux, 2-38
- .bash_profile file, 2-93
- Basic installation type
 - noninteractive installations, B-5
- block device
 - device name on HP-UX, 2-54, 2-71
 - device name on Linux, 2-57, 2-60
 - device name on Solaris, 2-64, 2-85, 2-88
 - device name on Tru64 UNIX, 2-65
- Bourne shell
 - default user startup file, 2-93
 - setting shell limits on Linux, 2-38
- bundle
 - checking on HP-UX, 2-10

C

- C compiler
 - requirement on AIX, 2-6
 - requirement on HP-UX, 2-9
 - requirement on Linux, 2-17
 - requirement on Tru64 UNIX, 2-23
 - See also* Pro*C/C++
- C shell
 - default user startup file, 2-93
 - setting shell limits on Linux, 2-38
- CD-ROMs
 - copying software from, 3-3
 - mounting
 - on AIX, A-2
 - on HP-UX, A-2
 - on Linux, A-3
 - on Solaris, A-4

- on Tru64 UNIX, A-5
- certification, hardware and software, 1-10
- cfgmgr command, 2-53, 2-67
- character device
 - device name on HP-UX, 2-54, 2-71
 - device name on Solaris, 2-64, 2-85, 2-88
 - device name on Tru64 UNIX, 2-65
- character raw device
 - device name on AIX, 2-53
- Character Set Scanner, 1-11
- character sets
 - AL24UTFFSS, 1-11
 - UTF8, 1-11
- chdev command, 2-67
- chmod command, 2-45, 2-48, 2-53, 2-54, 2-62, 2-63, 2-64, 2-66, 2-82
- chown command, 2-45, 2-48, 2-53, 2-54, 2-62, 2-63, 2-64, 2-66, 2-82
- CLASSPATH environment variable, 4-9
- cluster file system
 - storage option for datafiles, 2-46
- configuration assistants
 - troubleshooting, F-4
- control files
 - locating, 5-10
 - naming, D-7
 - raw devices for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - reviewing, 5-8
 - using Oracle Enterprise Manager Database Control with, 5-10
- CONTROL_FILES initialization parameter, 5-10
- corrective service diskette
 - See* CSD
- CSD
 - download location for MQSeries
 - on Tru64 UNIX, 2-25
 - download location for WebSphere MQ
 - on AIX, 2-8
 - on HP-UX, 2-12
 - on Linux, 2-18
 - on Solaris, 2-22
 - requirements
 - on AIX, 2-8
 - on HP-UX, 2-12
 - on Linux, 2-17
 - on Solaris, 2-22
 - on Tru64 UNIX, 2-25
- csh.login file, 2-38
- csh.login.local file, 2-38
- custom database
 - failure groups for ASM, 2-51
 - requirements when using ASM, 2-50
- Custom installation type
 - reasons for choosing, 2-26
- custom.rsp file, B-4

D

- data loss
 - minimizing with ASM, 2-51
- Database Configuration Assistant
 - troubleshooting, F-4
- Database Control
 - default ports, E-2
- databases
 - administrative information, D-3
 - ASM requirements, 2-50
 - files, 5-8
 - identifying, 5-7
 - initialization parameter file, 5-8
 - OFA file naming conventions, D-7
 - redo log files, 5-9
 - removing, 6-2
 - tablespaces, 5-8
- datafiles
 - creating separate directories for, 2-48
 - defined, 5-8
 - managing with ASM, 1-6
 - minimum disk space for, 2-47
 - naming, D-7
 - options for placing on file system, 2-46
 - recommendations for file system, 2-47
 - reviewing, 5-8
 - setting permissions on datafile directories, 2-48
 - setting up, 5-8
 - storage options, 2-46
- DB_DOMAIN initialization parameter, 5-7
- DB_NAME initialization parameter, 5-7
- dba group
 - and ASM disks on AIX, 2-53
 - and ASM disks on HP-UX, 2-54
 - and ASM disks on Linux, 2-62, 2-63, 2-82
 - and ASM disks on Solaris, 2-64
 - and ASM disks on Tru64 UNIX, 2-66
 - and SYSDBA privilege, 2-26
 - creating, 2-28
 - description, 2-26
 - granting privileges on HP-UX, 2-32
 - raw device group
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-82
 - on Solaris, 2-90
- DBCA_RAW_CONFIG environment variable, 2-95
 - on AIX, 2-71, 2-88
 - on HP-UX, 2-75
 - on Linux, 2-84
 - on Solaris, 2-92
- DBSNMP user password, 3-7
- DCE
 - patches required on HP-UX, 2-12
 - requirements for ASO on HP-UX, 2-9
 - requirements for ASO on Tru64 UNIX, 2-23
 - requirements on HP-UX, 2-9
 - requirements on Tru64 UNIX, 2-23
- default file mode creation mask
 - setting, 2-93
- device names
 - for raw devices on Linux, 2-82
 - IDE disks on Linux, 2-57, 2-60, 2-77, 2-79
 - on AIX, 2-53
 - on HP-UX, 2-54, 2-71
 - on Linux, 2-57, 2-60
 - on Solaris, 2-64, 2-85, 2-88
 - on Tru64 UNIX, 2-65
 - RAID on Linux, 2-57, 2-60, 2-77
 - raw devices on Linux, 2-62
 - SCSI disks on Linux, 2-57, 2-60, 2-77, 2-79
- device numbers
 - identifying major numbers on AIX, 2-67
 - identifying minor numbers on HP-UX, 2-72
- directory
 - creating separate datafile directories, 2-48
 - database file directory, 2-47
 - Oracle base directory, 2-42
 - Oracle home directory, 2-43
 - Oracle Inventory directory, 2-42
 - oraInventory, 2-42
 - permission for datafile directories, 2-48
- disk
 - changing permissions and owner for ASM
 - on HP-UX, 2-54
- disk devices
 - managing with ASM, 1-6
- disk group
 - ASM, 1-6, 2-50
 - recommendations for ASM disk groups, 2-50
- disk space
 - checking, 2-4
 - requirement for Oracle base directory, 2-44, 2-45
 - requirements for preconfigured database in ASM, 2-50
- disklabel command, 2-65
- disks
 - changing permissions and owner for ASM
 - on AIX, 2-53
 - on HP-UX, 2-54
 - on Solaris, 2-64
 - on Tru64 UNIX, 2-66
 - checking availability for ASM on AIX, 2-53, 2-57, 2-60, 2-77, 2-79
 - checking availability for ASM on HP-UX, 2-54
 - checking availability for ASM on Solaris, 2-64, 2-85
 - checking availability for ASM on Tru64 UNIX, 2-65
 - checking availability for raw devices on AIX, 2-67
 - checking availability for raw devices on HP-UX, 2-71
 - checking availability for raw devices on Solaris, 2-88
 - configuring for ASM on AIX, 2-53
 - configuring for ASM on HP-UX, 2-54
 - configuring for ASM on Linux, 2-55
 - configuring for ASM on Solaris, 2-63
 - configuring for ASM on Tru64 UNIX, 2-65
 - configuring new disks, 2-53, 2-67

- configuring raw devices on AIX, 2-67
- configuring raw devices on HP-UX, 2-71
- configuring raw devices on Solaris, 2-84
- device names on Solaris, 2-88
- displaying attached disks
 - on Linux, 2-57, 2-60, 2-77, 2-79
- identifying AdvFS file systems on Tru64 UNIX, 2-65
- identifying file systems on Solaris, 2-64, 2-86, 2-88
- identifying LSM disks on Tru64 UNIX, 2-66
- identifying LVM disks on AIX, 2-53, 2-67
- identifying LVM disks on HP-UX, 2-54, 2-72
- identifying LVM disks on Solaris, 2-63, 2-89
- identifying UFS file systems on Tru64 UNIX, 2-66
- identifying unused disks on Tru64 UNIX, 2-65
- DISPLAY environment variable, 2-35
 - setting, 2-93, 2-94
- Distributed Computing Environment
 - See DCE
- DOMAIN_NAME initialization parameter, 5-7
- drive failure safeguards, D-3
- dupatch command, 2-25
- Dynamic Server
 - See Informix Dynamic Server

E

- emulator
 - installing from X emulator, 2-2
- Enterprise Manager
 - default ports, E-2
- enterprise.rsp file, B-4
- env command, 2-96
- environment
 - checking settings, 2-96
 - configuring for oracle user, 2-93
- environment variables
 - DBCA_RAW_CONFIG, 2-71, 2-75, 2-84, 2-88, 2-92
 - DISPLAY, 2-35, 2-93, 2-94
 - ORACLE_BASE, 2-42, 2-45, 2-93
 - ORACLE_HOME, 2-93, 2-96
 - ORACLE_SID, 2-93
 - PATH, 2-93
 - removing from shell startup file, 2-94
 - SHELL, 2-93
 - TEMP and TMPDIR, 2-4, 2-95
 - TNS_ADMIN, 2-96
- errata
 - Linux kernel errata, 2-18
- errors
 - configuration assistants, F-4
 - installation, F-3, F-4
 - non-interactive installation, F-4
 - silent mode, F-4
 - X windows, F-2
- ESQL/C
 - See Informix ESQL/C
- /etc/csh.login file, 2-38
- /etc/csh.login.local file, 2-38

- /etc/oraInst.loc file, F-3
- /etc/pam.d/login file, 2-38
- /etc/privgroup file, 2-32
- /etc/profile file, 2-38
- /etc/profile.local file, 2-38
- /etc/raw file, 2-82
- /etc/security/limits.so file, 2-38
- /etc/sysconfig/rawdevices file, 2-62, 2-82
- /etc/sysconfigtab file, 2-41
- /etc/sysctl.conf file, 2-37
- /etc/system file, 2-39
- EXAMPLE tablespace, 5-9
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90

- examples
 - ASM failure groups, 2-51
 - Oracle base directories, 2-42
- exec_disable_arg_limit subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- external jobs
 - UNIX user required for, 2-26
- external redundancy
 - ASM redundancy level, 1-6
- extjob executable
 - changing permissions, 4-2
 - UNIX user required for, 2-26
- extjob user
 - creating on HP-UX, 2-31

F

- failure group
 - ASM, 2-50
 - characteristics of ASM failure group, 2-51
 - examples of ASM failure groups, 2-51
- fatal errors, F-4
- fdisk command, 2-57, 2-60, 2-77, 2-79
- fifo_do_adaptive subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- file command, 2-65
- file mode creation mask
 - setting, 2-93
- file naming, D-4
- file system
 - appropriate for Oracle base directory, 2-45
 - datafile and recovery file placement options, 2-46
 - identifying AdvFS file systems on Tru64 UNIX, 2-65
 - identifying formatted devices on Solaris, 2-64, 2-86, 2-88
 - identifying UFS file systems on Tru64 UNIX, 2-66
 - NFS, 1-11
 - organization, D-3
 - requirements for Oracle base directory, 2-45
 - storage option for datafiles, 2-46
 - using for datafiles, 2-47
 - writing to, 1-11

- file-max file, 2-36
- file-max parameter
 - recommended value on Linux, 2-36
- files, D-7
 - \$ORACLE_HOME/dbs/initsid.ora, 5-8
 - \$ORACLE_HOME/install/portlist.ini, 5-2
 - .bash_profile, 2-93
 - contol, D-7
 - control, 5-10
 - control files
 - raw devices for on AIX, 2-69
 - raw devices for on HP-UX, 2-73
 - raw devices for on Linux, 2-76
 - raw devices for on Solaris, 2-85, 2-90
 - custom.rsp, B-4
 - datafiles, D-7
 - default shell startup file, 2-38
 - editing shell startup file, 2-93
 - enterprise.rsp, B-4
 - /etc/csh.login, 2-38
 - /etc/csh.login.local, 2-38
 - /etc/group, D-5
 - /etc/oraInst.loc, F-3
 - /etc/pam.d/login, 2-38
 - /etc/passwd, D-5
 - /etc/privgroup, 2-32
 - /etc/profile, 2-38
 - /etc/profile.local, 2-38
 - /etc/raw, 2-82
 - /etc/security/limits.so, 2-38
 - /etc/sysconfig/rawdevices, 2-62, 2-82
 - /etc/sysconfigtab, 2-41
 - /etc/sysctl.conf, 2-37
 - /etc/system, 2-39
 - listener.ora, 4-5
 - .login, 2-93
 - mgw.ora, 4-9
 - oraInst.loc, 2-27, 2-43, B-3
 - oratab, 2-44
 - password file
 - raw device for on AIX, 2-69
 - raw device for on HP-UX, 2-73
 - raw device for on Linux, 2-76
 - raw device for on Solaris, 2-85, 2-90
 - /proc/sys/fs/file-max, 2-36
 - /proc/sys/kernel/sem, 2-36
 - /proc/sys/kernel/shmall, 2-36
 - /proc/sys/kernel/shmmax, 2-36
 - /proc/sys/kernel/shmmni, 2-36
 - /proc/sys/net/ipv4/ip_local_port_range, 2-36
 - .profile, 2-93
 - profile.local, 2-38
 - raw device mapping file, 2-95
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-83
 - on Solaris, 2-91
 - redo log, 5-9
 - redo log files
 - raw devices for on AIX, 2-69

- raw devices for on HP-UX, 2-73
- raw devices for on Linux, 2-76
- raw devices for on Solaris, 2-85, 2-90
- response files, B-4
- server parameter file
 - raw devices for on AIX, 2-69
 - raw devices for on HP-UX, 2-73
 - raw devices for on Linux, 2-76
 - raw devices for on Solaris, 2-85, 2-90
- SPFILE
 - raw device for on AIX, 2-69
 - raw device for on Solaris, 2-85
 - raw devices for on HP-UX, 2-73
 - raw devices for on Linux, 2-76
 - raw devices for on Solaris, 2-90
- standard.rsp, B-4
- tnsnames.ora, 4-4
 - /var/opt/oracle/oraInst.loc, F-3
- filesets
 - AIX requirements, 2-6
 - checking on AIX, 2-7
 - checking on HP-UX, 2-10
 - for raw devices on Linux, 2-82
- format command, 2-64, 2-85, 2-88

G

- Gateway
 - See* Oracle Messaging Gateway
- getconf command, 2-5
- getprivgrp command, 2-32
- gh_chunks subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- global database name, 5-7
- groupadd command, 2-28, 2-29
- groups
 - checking for existing oinstall group, 2-27
 - creating the dba group, 2-28
 - creating the oinstall group, 2-27
 - creating the oper group, 2-28
 - UNIX OSDBA group (dba), 2-26
 - UNIX OSOPER group (oper), 2-26

H

- hardware certification, 1-10
- hardware requirements, 2-3
- high redundancy
 - ASM redundancy level, 1-6
- home directory
 - distribution, D-3
 - login, D-3
 - See* Oracle home directory
- HP 9000 system
 - HP-UX operating system requirement, 2-9
- HP-UX
 - block and character device names, 2-54, 2-71
 - checking disk availability for ASM, 2-54
 - checking disk availability for raw devices, 2-71
 - checking patch requirements, 2-12

- checking quality pack requirement, 2-12
- checking software requirements, 2-10
- checking version, 2-10
- configuring disks for ASM, 2-54
- configuring kernel parameters, 2-34
- configuring raw devices, 2-71
- creating a raw device mapping file, 2-74
- creating a volume group, 2-71
- creating a volume group directory on HP-UX, 2-72
- creating an unprivileged user, 2-31
- creating logical volumes, 2-73
- creating raw logical volumes, 2-73
- creating required X library symbolic links, 2-13
- creating volume groups, 2-72
- granting privileges to OSDBA group, 2-32
- identifying disks for ASM, 2-54
- identifying disks for LVM, 2-71
- identifying LVM disks, 2-54, 2-72
- initializing disks for LVM, 2-72
- location of oratab file, 2-44
- maxvgs parameter, 2-72
- patch download location, 2-12
- patch requirements, 2-11
- quality pack download location, 2-12
- quality pack requirements, 2-11
- raw device sizes, 2-73
- requirements on HP 9000 systems, 2-9
- requirements on Itanium systems, 2-9
- setting kernel parameters, 2-35
- software requirements, 2-8
- specifying owner and permissions of raw devices, 2-74
- System Administration Manager (SAM), 2-35
- HP-UX Developer's Toolkit
 - Motif requirement, 2-9

I

- IBM MQSeries
 - requirement on Tru64 UNIX, 2-23
- IBM WebSphere MQ
 - requirement on AIX, 2-6
 - requirement on HP-UX, 2-9
 - requirement on Linux, 2-17
 - requirement on Solaris, 2-19
 - requirement on Tru64 UNIX, 2-23
- id command, 2-31
- IDE disks
 - device names on Linux, 2-57, 2-60, 2-77, 2-79
- Imake
 - requirement on HP-UX, 2-9
- Informix Dynamic Server
 - requirements on HP-UX, 2-10
 - requirements on Solaris, 2-20
- Informix ESQ/C
 - requirements on HP-UX, 2-10
 - requirements on Solaris, 2-20
- Ingres II
 - requirements on HP-UX, 2-10
- requirements on Solaris, 2-20
- initialization parameter file
 - description, 5-8
 - in databases, 5-8
 - initsid.ora, 5-8
- initialization parameters
 - DB_NAME, 5-7
 - DOMAIN_NAME, 5-7
 - SERVICE_NAMES, 5-7
- initsid.ora initialization parameter file, 5-8
- installation
 - overview, 3-2
 - available products, 1-3
 - completing, 3-7
 - considerations, 1-10
 - errors, F-3, F-4
 - silent mode, F-4
 - guidelines, 3-5
 - log files, F-3
 - non-interactive
 - error handling, F-5
 - oraInst.loc file, B-3
 - response files, B-4
 - preparing, B-4
 - silent mode, F-4
 - templates, B-4
 - silent mode, B-6, B-7
- installation guidelines, 3-5, 3-6
 - Oracle Label Security, 3-5
 - Oracle Universal Installer, 3-5
 - re-installing Oracle software, 3-6
- installation software
 - copying to a hard disk, 3-3
 - extracting, 3-3
 - re-installing, 3-6
- installation types
 - and ASM requirements, 2-50
- instance
 - instance identifier (SID), 2-93
- instfix command, 2-8
- I/O loads, D-4
 - OFA, D-3
- ioscan command, 2-54, 2-71
- ip_local_port_range file, 2-36
- ip_local_port_range parameter
 - recommended value on Linux, 2-36
- ipc
 - kernel subsystem on Tru64 UNIX, 2-40
- IPC protocol address
 - Oracle Messaging Gateway setting, 4-6
- isainfo command, 2-5
- iSQL*Plus
 - default ports, E-2
- Itanium
 - HP-UX operating system requirement, 2-9
- Itanium system
 - HP-UX operating system requirements, 2-9

J

Java

font package requirements for Solaris, 2-19

Java Development Kit

See JDK

Java Runtime Environment

See JRE

JDK

AIX requirements, 2-6

default locations, 3-7

font packages required on Solaris, 2-19

HP-UX requirements, 2-8

internationalization class, 4-9

runtime class, 4-9

Tru64 UNIX requirements, 2-23

JRE

Oracle supplied version, 3-7

K

kernel

Linux errata, 2-18

kernel parameters

checking on Linux, 2-37

checking on Solaris, 2-39

checking on Tru64 UNIX, 2-41

configuring on HP-UX, 2-34

configuring on Linux, 2-36

configuring on Solaris, 2-39

configuring on Tru64 UNIX, 2-40

making changes persist on Linux, 2-37

making changes persist on Solaris, 2-39

making changes persist on Tru64 UNIX, 2-41

maxvgs on HP-UX, 2-72

setting on HP-UX, 2-35

setting on Linux, 2-37

kernel subsystem attributes

See kernel parameters

Korn shell

default user startup file, 2-93

setting shell limits on Linux, 2-38

ksh

See Korn shell

ksi_alloc_max parameter

recommended value on HP-UX, 2-34

L

library

Motif library, 2-6

licensing information, 1-3

limit command, 2-38

limits.so file, 2-38

Linux

binding raw devices to partitions, 2-81

block device names, 2-57, 2-60

checking distribution, 2-18

checking existence of nobody user, 2-31

checking kernel parameters, 2-37

checking raw device bindings, 2-62

checking software requirements, 2-18

checking version, 2-18

configuring disks for ASM, 2-55

configuring kernel parameters, 2-36

creating a raw device mapping file, 2-83

creating partitions, 2-77

creating raw partitions, 2-76

displaying attached disks, 2-57, 2-60, 2-77, 2-79

IDE disk device names, 2-57, 2-60, 2-77, 2-79

identifying disks for ASM, 2-57, 2-60, 2-77, 2-79

kernel errata, 2-18

location of oratab file, 2-44

making kernel parameter changes persist, 2-37

RAID device names, 2-57, 2-60, 2-77

raw device sizes, 2-76

raw devices, 2-62

SCSI disk device names, 2-57, 2-60, 2-77, 2-79

setting kernel parameters, 2-37

setting shell limits, 2-38

software requirements, 2-14

specifying owner and permissions on raw devices, 2-82

listener

identifying Oracle home for, 2-92

lsnrctl command, 2-93

stopping, 2-92, 2-93

stopping existing listener process, 2-92

listener.ora file, 4-5

modifying for external procedures, 4-6

on AIX, 4-6

on HP-UX, 4-6

on Linux, 4-6

on Tru64 UNIX, 4-6

local device

using for datafiles, 2-47

log files, F-3

troubleshooting, F-3

logical volume manager

See LVM

logical volumes

creating on AIX, 2-69

creating on HP-UX, 2-73

creating on Solaris, 2-90

.login file, 2-93

login file, 2-38

lsdev command, 2-53, 2-57, 2-60, 2-67, 2-77, 2-79

lspp command, 2-7, 2-8

LSM

identifying volume group devices on Tru64 UNIX, 2-66

lsnrctl command, 2-93

lspv command, 2-53, 2-67

lvcreate command, 2-73

LVM

checking disk availability on HP-UX, 2-71

configuring raw devices on AIX, 2-67

configuring raw devices on HP-UX, 2-71

configuring raw devices on Solaris, 2-84

creating a volume group directory on HP-UX, 2-72

- creating a volume group on AIX, 2-67
- creating a volume group on HP-UX, 2-71
- creating a volume group on Solaris, 2-88
- creating raw logical volumes on AIX, 2-68
- creating raw logical volumes on HP-UX, 2-73
- creating raw logical volumes on Solaris, 2-89
- creating volume groups on AIX, 2-68
- creating volume groups on HP-UX, 2-72
- creating volume groups on Solaris, 2-89
- identifying available disks on AIX, 2-67
- identifying available disks on HP-UX, 2-71
- identifying available disks on Solaris, 2-88
- identifying LVM disks on HP-UX, 2-72
- identifying major device numbers on AIX, 2-67
- identifying minor device numbers on HP-UX, 2-72
- identifying volume group devices on AIX, 2-53, 2-67
- identifying volume group devices on HP-UX, 2-54
- identifying volume group devices on Solaris, 2-63, 2-89
- initializing disks on AIX, 2-67
- initializing disks on HP-UX, 2-72
- initializing disks on Solaris, 2-89
- maximum number of volume groups on HP-UX, 2-72
- recommendations for ASM, 2-50
- Veritas Volume Manager on Solaris, 2-63, 2-89

M

- maintenance level
 - checking on AIX, 2-7
- major device numbers
 - identifying on AIX, 2-67
- mapping file
 - for raw devices, 2-95
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-83
 - on Solaris, 2-91
- mask
 - setting default file mode creation mask, 2-93
- max_async_req subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- max_objvs subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- max_per_proc_address_space subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- max_per_proc_data_size subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- max_per_proc_stack_size subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- max_sessions subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- max_thread_proc parameter
 - recommended value on HP-UX, 2-34
- maxdsiz parameter
 - recommended value on HP-UX, 2-34

- maxdsiz_64bit parameter
 - recommended value on HP-UX, 2-34
- maxssiz parameter
 - recommended value on HP-UX, 2-34
- maxssiz_64bit parameter
 - recommended value on HP-UX, 2-34
- maxswapchunks parameter
 - recommended value on HP-UX, 2-34
- maxuprc parameter
 - recommended value on HP-UX, 2-34
- maxvgs parameter, 2-72
- memory requirements, 2-3
- Messaging Gateway
 - See Oracle Messaging Gateway
- MGMT_ECM_DEPOT_TS tablespace, 5-9
- MGMT_TABLESPACE tablespace, 5-9
- MGW_AGENT service name, 4-9
- mgwextproc service
 - adding static service information, 4-6
- mgw.ora file
 - modifying, 4-9
- migrating
 - See upgrading
- minor device numbers
 - identifying on HP-UX, 2-72
- mkdir command, 2-45, 2-48
- mklv command, 2-69
- mknod command, 2-72
- mkvq command, 2-68
- mode
 - setting default file mode creation mask, 2-93
- Motif
 - requirement on Tru64 UNIX, 2-23
- mount command, 2-66
- mount point
 - for Oracle base directory, 2-42
- mount point directories, 3-4
- MQSeries
 - checking requirements on Tru64 UNIX, 2-24
 - class, 4-9
 - CSD download location
 - for Tru64 UNIX, 2-25
 - CSDs required
 - on Tru64 UNIX, 2-25
 - requirement on Tru64 UNIX, 2-23
- msg_size subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- msgmap parameter
 - recommended value on HP-UX, 2-34
- msgmni parameter
 - recommended value on HP-UX, 2-34
- msgseg parameter
 - recommended value on HP-UX, 2-34
- msgtql parameter
 - recommended value on HP-UX, 2-34
- multiple Oracle homes, 1-10

N

- naming subdirectories, D-6

- NAS
 - usage guidelines, C-1
 - using with ASM, C-5
 - NCR Teradata ODBC Driver
 - requirements on HP-UX, 2-10
 - requirements on Solaris, 2-20
 - ncsize parameter
 - recommended value on HP-UX, 2-34, 2-35
 - Net Configuration Assistant
 - troubleshooting, F-4
 - Network File System *See* NFS
 - Network Information Services
 - See* NIS
 - new_wire_method subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
 - nfile parameter
 - recommended value on HP-UX, 2-35
 - nflocks parameter
 - recommended value on HP-UX, 2-35
 - NFS
 - usage guidelines, C-1
 - using for installation, 1-11
 - using with ASM, C-5
 - ninode parameter
 - recommended value on HP-UX, 2-35
 - NIS
 - alternative to local users and groups, 2-27
 - nkthread parameter
 - recommended value on HP-UX, 2-35
 - nobody user
 - checking existence of, 2-31
 - description, 2-26
 - nofile
 - shell limit on Linux, 2-38
 - non-interactive installation
 - oraInst.loc file, B-3
 - response files
 - preparing, B-4
 - templates, B-4
 - silent mode, B-6, B-7
 - errors, F-4
 - non-interactive installations
 - running
 - Oracle Universal Installer, B-6
 - normal redundancy
 - ASM redundancy level, 1-6
 - nproc
 - shell limit on Linux, 2-38
 - nproc parameter
 - recommended value on HP-UX, 2-35
- O**
-
- ODBC driver
 - Teradata requirement on HP-UX, 2-10
 - Teradata requirement on Solaris, 2-20
 - OEM
 - See* Oracle Enterprise Manager
 - OFA
 - administrative information, D-3
 - ASM, D-7
 - concurrent execution of application software, D-3
 - distributed I/O loads, D-3
 - drive failure safeguards, D-3
 - file identification, D-8
 - file mapping, D-3, D-9
 - file naming, D-4
 - file system organization, D-3
 - files systems, D-4
 - hardware support, D-3
 - home directory distribution, D-3
 - I/O loads, D-4
 - login home directories, D-3
 - naming, D-4
 - database files, D-7
 - Oracle base directory, D-5
 - subdirectories, D-6
 - tablespaces, D-8
 - very large databases, D-4
 - on UNIX, D-4
 - Oracle Managed Files, D-7
 - pathnames, D-5
 - recommendations for Oracle base directory, 2-42
 - recommended path for Oracle base
 - directory, 2-42
 - recommended path for Oracle home
 - directory, 2-43
 - recommended path for Oracle Inventory
 - directory, 2-42
 - special tablespaces, D-8
 - standard, D-1
 - tablespace contents, D-4
 - UNIX directory subtrees, D-3
 - using separate segments, D-7
 - oinstall group
 - checking for existing, 2-27
 - creating, 2-27
 - description, 2-26
 - OLS
 - See* Oracle Label Security
 - OMF
 - See* Oracle Managed Files
 - Open Client
 - See* Sybase Open Client
 - Open Database Connectivity
 - See* ODBC
 - oper group
 - and SYSOPER privilege, 2-26
 - creating, 2-28, 2-29
 - description, 2-26
 - operating system
 - checking distribution and version of Linux, 2-18
 - checking version of AIX, 2-7
 - checking version of HP-UX, 2-10
 - checking version of Solaris, 2-20
 - checking version of Tru64 UNIX, 2-24
 - operating system requirements
 - AIX, 2-6
 - HP-UX, 2-8
 - Linux, 2-14

- Solaris, 2-19
- Tru64 UNIX, 2-23
- Optimal Flexible Architecture
 - See* OFA
- Oracle Advanced Security
 - patches required on HP-UX, 2-12
 - requirements on HP-UX, 2-9
 - requirements on Tru64 UNIX, 2-23
- Oracle base directory
 - and ORACLE_BASE environment variable, 2-42
 - creating, 2-45
 - creating new, 2-45
 - description, 2-42
 - determining disk space on, 2-44
 - disk space requirements, 2-44, 2-45
 - equivalent directory on Windows, 2-42
 - examples, 2-42
 - identifying appropriate file system, 2-45
 - identifying existing, 2-43
 - mount point for, 2-42
 - naming conventions, D-5
 - recommended path, 2-42
 - relationship with Oracle Software Owner user, 2-42
 - requirement for, 2-42
 - requirements for existing directory, 2-44
 - requirements on file system, 2-45
- Oracle Cluster Registry
 - See* OCR
- Oracle Database
 - creating datafile directories, 2-48
 - Custom installation, 1-3
 - datafile storage options, 2-46
 - Enterprise Edition installation, 1-3
 - minimum disk space requirements, 2-47
 - privileged groups, 2-26
 - requirements with ASM, 2-50
 - setting ORACLE_SID environment variable, 2-93
 - Standard Edition installation, 1-3
- Oracle DCE Integration
 - requirements on HP-UX, 2-9
 - requirements on Tru64 UNIX, 2-23
- Oracle Enterprise Manager, 1-6
 - Database Control
 - logging into, 5-2
 - port number, 5-2
 - using to modify control files, 5-10
 - using to modify redo log files, 5-10
 - using to view control files, 5-10
 - using to view redo log files, 5-10
 - login privileges, 5-2
- Oracle home directory
 - description, 2-43
 - identifying for listener, 2-92
 - recommended path, 2-43
 - requirement for, 2-43
 - requirements, 2-43
 - using to identify Oracle base directory, 2-44
- Oracle home name, 2-43
- Oracle internationalization class, 4-9
- Oracle Inventory
 - description, 2-42
 - pointer file, 2-27
- Oracle Inventory directory
 - description, 2-42
 - recommended path, 2-42
- Oracle Inventory group
 - checking for existing, 2-27
 - creating, 2-27, 2-28
 - description, 2-26
- Oracle JDBC class, 4-9
- Oracle Label Security
 - installation guidelines, 3-5
 - post-installation tasks, 4-5
- Oracle Managed Files
 - OFA naming conventions, D-7
- Oracle Messaging Gateway
 - checking requirements on Tru64 UNIX, 2-24
 - CSD requirements
 - on AIX, 2-8
 - on HP-UX, 2-12
 - on Linux, 2-17
 - on Solaris, 2-22
 - on Tru64 UNIX, 2-25
 - requirements on AIX, 2-6
 - requirements on HP-UX, 2-9
 - requirements on Linux, 2-17
 - requirements on Solaris, 2-19
 - requirements on Tru64 UNIX, 2-23
- Oracle Messaging Gateway class, 4-9
- Oracle Net
 - configuration file directory, 4-4
 - identifying Oracle home for listener, 2-92
 - lsnrctl command, 2-93
 - stopping existing listener, 2-92
 - stopping listener, 2-92
 - stopping the listener, 2-93
- Oracle Net Services
 - post-installation tasks, 4-4
- Oracle Precompilers
 - post-installation tasks, 4-10
- Oracle software
 - removing, 6-7
- Oracle Software Owner user
 - and ASM disks, 2-53, 2-54, 2-62, 2-63, 2-64, 2-66, 2-82
 - configuring environment for, 2-93
 - creating, 2-29
 - description, 2-26
 - determining default shell, 2-93
 - raw device owner
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-82
 - on Solaris, 2-90
 - relationship with Oracle base directory, 2-42
 - required group membership, 2-26
 - setting shell limits for on Linux, 2-38
- Oracle Spatial
 - Motif requirement on AIX, 2-6

- Motif requirement on HP-UX, 2-9
 - Motif requirement on Solaris, 2-19
 - Motif requirement on Tru64 UNIX, 2-23
 - requirement for example programs
 - on AIX, 2-6
 - on HP-UX, 2-9
 - on Linux, 2-17
 - on Solaris, 2-19
 - on Tru64 UNIX, 2-23
 - X window requirement on Linux, 2-17
 - Oracle Transparent Gateway
 - Informix requirements on HP-UX, 2-10
 - Informix requirements on Solaris, 2-20
 - Ingres requirements on HP-UX, 2-10
 - Ingres requirements on Solaris, 2-20
 - Sybase requirements on AIX, 2-6
 - Sybase requirements on HP-UX, 2-10
 - Sybase requirements on Solaris, 2-20
 - Sybase requirements on Tru64 UNIX, 2-24
 - Teradata requirements on HP-UX, 2-10
 - Teradata requirements on Solaris, 2-20
 - Oracle Universal Installer
 - installation guidelines, 3-5
 - response files
 - list of, B-4
 - running, 3-6
 - oracle user
 - and ASM disks, 2-53, 2-54, 2-62, 2-63, 2-64, 2-66, 2-82
 - configuring environment for, 2-93
 - creating, 2-29, 2-30
 - description, 2-26
 - determining default shell, 2-93
 - raw device owner
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-82, 2-90
 - relationship with Oracle base directory, 2-42
 - required group membership, 2-26
 - setting shell limits for on Linux, 2-38
 - ORACLE_BASE environment variable, 2-42, 2-45
 - removing from shell startup file, 2-94
 - setting, 2-93
 - ORACLE_HOME environment variable
 - removing from shell startup file, 2-94
 - setting, 2-93
 - unsetting, 2-96
 - ORACLE_SID environment variable
 - removing from shell startup file, 2-94
 - setting, 2-93
 - oraInst.loc file, 2-43
 - location, 2-27
 - location of, 2-27
 - oraInventory directory
 - See Oracle Inventory directory
 - oratab file
 - format, 2-44
 - location of, 2-44
 - OSDBA group
 - and ASM disks on AIX, 2-53
 - and ASM disks on HP-UX, 2-54
 - and ASM disks on Linux, 2-62, 2-63, 2-82
 - and ASM disks on Solaris, 2-64
 - and ASM disks on Tru64 UNIX, 2-66
 - and SYSDBA privilege, 2-26
 - creating, 2-28
 - description, 2-26
 - granting privileges on HP-UX, 2-32
 - raw device group
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-82
 - on Solaris, 2-90
 - oslevel command, 2-7
 - OSOPER group
 - and SYSOPER privilege, 2-26
 - creating, 2-28
 - description, 2-26
 - OTN Web site
 - downloading installation software from, 3-2
 - OUI
 - See Oracle Universal Installer
- ## P
-
- packages
 - checking on Linux, 2-18
 - checking on Solaris, 2-20
 - requirements for Solaris, 2-19
 - partition
 - using with ASM, 2-50
 - partitions
 - binding to raw devices on Linux, 2-81
 - creating on Linux, 2-77
 - creating raw partitions on Linux, 2-76
 - required sizes for raw devices on Linux, 2-76
 - slices on Solaris, 2-64, 2-86, 2-88
 - passwd command, 2-30
 - password file
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - passwords
 - DBSNMP, 3-7
 - guidelines, 3-7
 - resetting, 5-6
 - with Database Control, 5-6
 - with SQL*Plus, 5-6
 - reviewing, 5-3
 - SYS, 3-7
 - SYSMAN, 3-7
 - SYSTEM, 3-7
 - unlocking, 5-6
 - with Database Control, 5-6
 - with SQL*Plus, 5-6
 - patchadd command, 2-22
 - patches
 - checking on AIX, 2-8

- checking on HP-UX, 2-12
- checking on Solaris, 2-22
- download location for AIX, 2-8
- download location for HP-UX, 2-12
- download location for Solaris, 2-22
- download location for Tru64 UNIX, 2-25
- requirements on AIX, 2-7
- requirements on HP-UX, 2-11
- requirements on Solaris, 2-21
- requirements on Tru64 UNIX, 2-25
- PATH environment variable
 - setting, 2-93
- pathnames
 - OFA, D-5
- PC X server
 - installing from, 2-2
- per_proc_address subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- per_proc_data_size subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- per_proc_stack_size subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- permissions
 - for datafile directories, 2-48
 - for Oracle base directory, 2-45
- physical RAM requirements, 2-3
- pkginfo command, 2-20
- ports
 - default port ranges, E-1
- post-installation
 - recommended tasks
 - root.sh script, backing up, 4-3
 - user accounts, setting up, 4-3
 - required tasks, 4-2
 - extjob executable, changing permissions, 4-2
 - Oracle Label Security, configuring, 4-5
 - Oracle Net Services, configuring, 4-4
 - Oracle Precompilers, 4-10
 - patches, installing and downloading, 4-2
- Precompilers
 - requirements on AIX, 2-6
 - requirements on HP-UX, 2-9
 - requirements on Linux, 2-17
 - requirements on Tru64 UNIX, 2-23
- preconfigured database
 - ASM disk space requirements, 2-50
 - requirements when using ASM, 2-50
- privgroup file, 2-32
- privileged groups
 - for Oracle Database, 2-26
- Pro*C/C++
 - configuring, 4-10
 - patches required on HP-UX, 2-12
 - patches required on Solaris, 2-22
 - PTFs and APARs required on AIX, 2-7
 - requirements on AIX, 2-6
 - requirements on HP-UX, 2-9
 - requirements on Linux, 2-17
 - requirements on Tru64 UNIX, 2-23
 - See also* C compiler

- proc
 - kernel subsystem on Tru64 UNIX, 2-40
- process
 - stopping existing, 2-92
 - stopping existing listener process, 2-92
 - stopping listener process, 2-92
- processor
 - checking system architecture, 2-5
 - /proc/sys/fs/file-max file, 2-36
 - /proc/sys/kernel/sem file, 2-36
 - /proc/sys/kernel/shmall file, 2-36
 - /proc/sys/kernel/shmmni file, 2-36
 - /proc/sys/net/ipv4/ip_local_port_range file, 2-36
- product
 - checking on HP-UX, 2-10
- .profile file, 2-93
- profile file, 2-38
- profile.local file, 2-38
- program technical fix
 - See* PTF
- ps command, 2-92
- PTF
 - checking, 2-8
 - download location, 2-8
 - requirements on AIX, 2-7
- pvcreate command, 2-72
- pvdisplay command, 2-54, 2-72

Q

- quality pack
 - checking requirements on HP-UX, 2-12
 - download location for HP-UX, 2-12
 - requirements on HP-UX, 2-11

R

- RAC
 - configuring disks for ASM on Linux, 2-78, 2-79
 - configuring disks for ASM on Solaris, 2-85
- rad_gh_regions attribute
 - recommended value on Tru64 UNIX, 2-40
- RAID
 - device names on Linux, 2-57, 2-60, 2-77
 - recommended ASM redundancy level, 2-50
 - using for Oracle datafiles, 2-47
- RAM requirements, 2-3
- raw command, 2-62, 2-81, 2-82
- raw device
 - for password file
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - for server parameter file
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - for SPFILE

- on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
- raw devices
 - binding on Linux, 2-62
 - binding to partitions on Linux, 2-81
 - block and character device names on
 - HP-UX, 2-71
 - block and character device names on Solaris, 2-88
 - changing permissions and owner for ASM on
 - Linux, 2-62, 2-63, 2-82
 - checking availability of disks on AIX, 2-67
 - checking disk availability on HP-UX, 2-71
 - checking disk availability on Solaris, 2-88
 - configuring on AIX, 2-67
 - configuring on HP-UX, 2-71
 - configuring on Solaris, 2-84
 - creating a raw device mapping file
 - on AIX, 2-70
 - on HP-UX, 2-74
 - on Linux, 2-83
 - on Solaris, 2-91
 - creating logical volumes on AIX, 2-69
 - creating logical volumes on HP-UX, 2-73
 - creating logical volumes on Solaris, 2-90
 - creating partitions on Linux, 2-77
 - creating raw logical volumes on AIX, 2-68
 - creating raw logical volumes on HP-UX, 2-73
 - creating raw logical volumes on Solaris, 2-89
 - creating raw partitions on Linux, 2-76
 - DBCA_RAW_CONFIG environment
 - variable, 2-95
 - device names on Linux, 2-62, 2-82
 - device names, 2-82
 - for control files
 - on AIX, 2-69, 2-85
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-90
 - for EXAMPLE tablespace
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - for redo log files
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - for SYSAUX tablespace
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-84, 2-89
 - for SYSTEM tablespace
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-84, 2-89
 - for TEMP tablespace
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - for UNDOTBS tablespace
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-84, 2-90
 - for USER tablespace
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - identifying disks on AIX, 2-67
 - identifying disks on HP-UX, 2-71
 - identifying disks on Solaris, 2-88
 - identifying LVM disks on HP-UX, 2-72
 - initializing disks for LVM on AIX, 2-67
 - initializing disks for LVM on HP-UX, 2-72
 - initializing disks for LVM on Solaris, 2-89
 - required sizes
 - on AIX, 2-68
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-89
 - specifying location of mapping file, 2-95
 - specifying owner and permissions on AIX, 2-70
 - specifying owner and permissions on
 - HP-UX, 2-74
 - specifying owner and permissions on Linux, 2-82
 - specifying owner and permissions on
 - Solaris, 2-90
 - storage option for datafiles, 2-46
 - value for DBCA_RAW_CONFIG environment
 - variable
 - on AIX, 2-71, 2-88
 - on HP-UX, 2-75
 - on Linux, 2-84
 - on Solaris, 2-92
- raw file, 2-82
- raw logical volumes
 - See* raw devices
- raw partitions
 - configuring on Solaris, 2-84
 - See* raw devices
- rawdevices file, 2-62, 2-82
- rdg
 - kernel subsystem on Tru64 UNIX, 2-40
- rdg_auto_msg_wires subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- rdg_max_auto_msg_wires subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- reboot command, 2-39
- recovery files
 - options for placing on file system, 2-46
- Red Hat
 - binding partitions to raw devices, 2-81
 - operating system requirements, 2-14, 2-15, 2-17

- Red Hat Package Manager
 - See RPM
- redo log, D-7
- redo log files
 - in starter database, 5-9
 - locating, 5-9
 - naming, D-7
 - raw devices for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
 - reviewing, 5-8
 - using Oracle Enterprise Manager Database Control
 - with, 5-10
- redundancy level
 - and space requirements for preconfigured
 - database, 2-50
 - for ASM, 1-6
- redundant array of independant disks
 - See RAID
- Rendezvous
 - requirement on HP-UX, 2-9
 - requirement on Linux, 2-17
 - requirement on Solaris, 2-19
- requirements
 - hardware, 2-3
- response files
 - custom.rsp, B-4
 - enterprise.rsp, B-4
 - standard.rsp, B-4
- root user, 3-6
 - logging in as, 2-2
- rootpre.sh script, 3-5
- root.sh script
 - backing up, 4-3
- RPM
 - checking on Linux, 2-18
- rpm command, 2-18
- rt
 - kernel subsystem on Tru64 UNIX, 2-40

S

- SAM
 - starting on HP-UX, 2-35
- sam command, 2-35
- SCSI disks
 - device names on Linux, 2-57, 2-60, 2-77, 2-79
- sem file, 2-36
- seminfo_semmni parameter
 - recommended value on Solaris, 2-39
- seminfo_semmns parameter
 - recommended value on Solaris, 2-39
- seminfo_semmsl parameter
 - recommended value on Solaris, 2-39
- seminfo_semvmx parameter
 - recommended value on Solaris, 2-39
- semmap parameter
 - recommended value on HP-UX, 2-35
- semmni parameter
 - recommended value on HP-UX, 2-35
 - recommended value on Linux, 2-36
 - recommended value on Solaris, 2-39
- semmns parameter
 - recommended value on HP-UX, 2-35
 - recommended value on Linux, 2-36
 - recommended value on Solaris, 2-39
- semmnu parameter
 - recommended value on HP-UX, 2-35
- semmsl parameter
 - recommended value on Linux, 2-36
 - recommended value on Solaris, 2-39
- semopm parameter
 - recommended value on Linux, 2-36
- semvmx parameter
 - recommended value on HP-UX, 2-35
 - recommended value on Solaris, 2-39
- server parameter file
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
- service command, 2-62, 2-82
- SERVICE_NAMES initialization parameter, 5-7
- setld command, 2-24
- setprivgrp command, 2-32
- shell
 - default shell startup file, 2-38
 - determining default shell for oracle user, 2-93
- SHELL environment variable
 - checking value of, 2-93
- shell limits
 - setting on Linux, 2-38
- shell startup file
 - editing, 2-93
 - removing environment variables, 2-94
- shm_max subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- shm_min subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- shm_mni subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- shm_seg subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- shmall file, 2-36
- shmall parameter
 - recommended value on Linux, 2-36
- shminfo_shmmax parameter
 - recommended value on Solaris, 2-39
- shminfo_shmmin parameter
 - recommended value on Solaris, 2-39
- shminfo_shmmni parameter
 - recommended value on Solaris, 2-39
- shminfo_shmseg parameter
 - recommended value on Solaris, 2-39
- shmmax file, 2-36
- shmmax parameter
 - recommended value on HP-UX, 2-35

- recommended value on Linux, 2-36
- recommended value on Solaris, 2-39
- shmmin parameter
 - recommended value on Solaris, 2-39
- shmmni file, 2-36
- shmmni parameter
 - recommended value on HP-UX, 2-35
 - recommended value on Linux, 2-36
 - recommended value on Solaris, 2-39
- shmseg parameter
 - recommended value on HP-UX, 2-35
 - recommended value on Solaris, 2-39
- shutdown command, 2-41
- SID, 5-7
 - setting ORACLE_SID environment variable, 2-93
- silent mode installation, B-6, B-7
- sizer command, 2-24
- smit command, 2-28, 2-29, 2-30
- software certification, 1-10
- software requirements
 - AIX, 2-6
 - checking on AIX, 2-7
 - checking on HP-UX, 2-10
 - checking on Linux, 2-18
 - checking on Solaris, 2-20
 - checking on Tru64 UNIX, 2-24
 - HP-UX, 2-8
 - Linux, 2-14
 - Solaris, 2-19
 - Tru64 UNIX, 2-23
- Solaris
 - block and character device names, 2-64, 2-85, 2-88
 - checking disk availability for ASM, 2-64, 2-85
 - checking disk availability for raw devices, 2-88
 - checking existence of nobody user, 2-31
 - checking kernel parameters, 2-39
 - checking patch requirements, 2-22
 - checking software requirements, 2-20
 - checking version, 2-20
 - configuring disks for ASM, 2-63
 - configuring kernel parameters, 2-39
 - configuring raw devices, 2-84
 - configuring raw partitions, 2-84
 - creating a raw device mapping file, 2-91
 - creating a volume group, 2-88
 - creating logical volumes, 2-90
 - creating raw logical volumes, 2-89
 - creating volume groups, 2-89
 - font packages for Java, 2-19
 - identifying disks for ASM, 2-64, 2-85
 - identifying disks for LVM, 2-88
 - identifying file systems, 2-64, 2-86, 2-88
 - identifying LVM disks, 2-63, 2-89
 - initializing disks for LVM, 2-89
 - location of oratab file, 2-44
 - making kernel parameter changes persist, 2-39
 - operating system requirements, 2-19
 - package requirements, 2-19
 - patch download location, 2-22
 - patch requirements, 2-21
 - raw device sizes, 2-89
 - software requirements, 2-19
 - specifying owner and permissions on raw devices, 2-90
 - Veritas Volume Manager, 2-63, 2-89
- SPFILE
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
- SQLJ class, 4-9
- ssm_threshold subsystem attribute
 - recommended value on Tru64 UNIX, 2-40
- standard.rsp files, B-4
- startup file
 - default shell startup file, 2-38
 - for shell, 2-93
- static service information
 - adding for mgwextproc service, 4-6
- subsets
 - checking on Tru64 UNIX, 2-24
 - requirements for Tru64 UNIX, 2-23
- SuSE
 - binding partitions to raw devices, 2-82
 - operating system requirements, 2-14, 2-15, 2-17
- swap space
 - requirements, 2-3
- swlist command, 2-10, 2-12
- Sybase Adaptive Server Enterprise
 - requirements on AIX, 2-6
 - requirements on HP-UX, 2-10
 - requirements on Solaris, 2-20
 - requirements on Tru64 UNIX, 2-24
- Sybase Open Client
 - requirements on AIX, 2-6
 - requirements on HP-UX, 2-10
 - requirements on Solaris, 2-20
 - requirements on Tru64 UNIX, 2-24
- symbolic links
 - X library links required on HP-UX, 2-13
- SYS user password, 3-7
- SYSAUX tablespace, 5-9
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-84, 2-89
- sysconfig command, 2-41
- sysconfigdb command, 2-41
- sysconfigtab file, 2-41
- sysctl command, 2-37
- sysctl.conf file, 2-37
- SYSDBA privilege
 - associated UNIX group, 2-26
- sysdef command, 2-39
- SYSMAN user password, 3-7
- SYSOPER privilege
 - associated UNIX group, 2-26
- System Administration Manager

- See SAM
- system architecture
 - checking, 2-5
- system file, 2-39
- System Identifier, 5-7
 - See SID
- SYSTEM tablespace, 5-9
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-84, 2-89
- SYSTEM user password, 3-7

T

- tablespaces
 - defined, 5-8
 - in databases, 5-8
 - reviewing, 5-8
 - setting up, 5-8
 - EXAMPLE, 5-9
 - MGMT_ECM_DEPOT_TS, 5-9
 - MGMT_TABLESPACE, 5-9
 - OFA
 - naming, D-8
 - special tablespaces, D-8
 - SYSAUX, 5-9
 - SYSTEM, 5-9
 - TEMP, 5-9
 - UNDOTBS, 5-9
 - USERS, 5-9
- tcsh shell
 - setting shell limits on Linux, 2-38
- TEMP environment variable, 2-4
 - setting, 2-95
- TEMP tablespace, 5-9
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
- temporary directory, 2-4
- temporary disk space
 - checking, 2-4
 - freeing, 2-4
 - requirements, 2-3
- Teradata ODBC Driver
 - requirements on HP-UX, 2-10
 - requirements on Solaris, 2-20
- TIBCO Rendezvous
 - requirement on HP-UX, 2-9
 - requirement on Linux, 2-17
 - requirement on Solaris, 2-19
- /tmp directory
 - checking space in, 2-4
 - freeing space in, 2-4
- TMPDIR environment variable, 2-4
 - setting, 2-95
- TNS_ADMIN environment variable

- unsetting, 2-96
- tnsnames.ora file, 4-4
 - adding a connect descriptor, 4-9
 - MGW_AGENT service name, 4-9
 - modifying for external procedures, 4-9
- Transparent Gateway
 - See Oracle Transparent Gateway
- troubleshooting, F-1
 - fatal errors, F-4
- Tru64 UNIX
 - block and character device names, 2-65
 - checking disk availability for ASM, 2-65
 - checking existence of nobody user, 2-31
 - checking kernel subsystem attributes, 2-41
 - checking software requirements, 2-24
 - checking version, 2-24
 - configuring disks for ASM, 2-65
 - configuring kernel parameters, 2-40
 - identifying AdvFS file systems, 2-65
 - identifying disks for ASM, 2-65
 - identifying LSM disks, 2-66
 - identifying UFS file systems, 2-66
 - identifying unused disks, 2-65
 - location of oratab file, 2-44
 - making subsystem attribute value changes
 - persist, 2-41
 - patch download location, 2-25
 - patch requirements, 2-25
 - software requirements, 2-23
 - subset requirements, 2-23

U

- UFS
 - identifying UFS file systems on Tru64 UNIX, 2-66
- ulimit command, 2-38
- Ultra Search
 - default ports, E-2
- umask, 2-96
- umask command, 2-93, 2-96
- uname command, 2-10, 2-20
- UNDOTBS tablespace, 5-9
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-84, 2-90
- UNIX commands, 2-96
 - cfgmgr, 2-53, 2-67
 - chdev, 2-67
 - chmod, 2-45, 2-48, 2-53, 2-54, 2-62, 2-63, 2-64, 2-66, 2-82
 - chown, 2-45, 2-48, 2-53, 2-54, 2-62, 2-63, 2-64, 2-66, 2-82
 - disklabel, 2-65
 - dupatch, 2-25
 - env, 2-96
 - fdisk, 2-57, 2-60, 2-77, 2-79
 - file, 2-65
 - format, 2-64, 2-85, 2-88

- getconf, 2-5
- getprivgrp, 2-32
- groupadd, 2-28, 2-29
- id, 2-31
- instfix, 2-8
- ioscan, 2-54, 2-71
- isainfo, 2-5
- limit, 2-38
- lsdev, 2-53, 2-57, 2-60, 2-67, 2-77, 2-79
- lslpp, 2-7, 2-8
- lspv, 2-53, 2-67
- lvcreate, 2-73
- mkdir, 2-45, 2-48
- mklv, 2-69
- mknod, 2-72
- mkvg, 2-68
- mount, 2-66
- oslevel, 2-7
- passwd, 2-30
- patchadd, 2-22
- pkginfo, 2-20
- ps, 2-92
- pvcreeate, 2-72
- pvdisplay, 2-54, 2-72
- raw, 2-62, 2-81, 2-82
- reboot, 2-39
- rpm, 2-18
- sam, 2-35
- service, 2-62, 2-82
- setld, 2-24
- setprivgrp, 2-32
- shutdown, 2-41
- sizer, 2-24
- smit, 2-28, 2-29, 2-30
- swap, 2-4
- swapon, 2-4
- swlist, 2-10, 2-12
- sysconfig, 2-41
- sysconfigdb, 2-41
- sysctl, 2-37
- sysdef, 2-39
- ulimit, 2-38
- umask, 2-93
- uname, 2-10, 2-20
- unset, 2-96
- unsetenv, 2-96
- useradd, 2-30
- vgchange, 2-72
- vgcreate, 2-72
- voldisk, 2-66
- vxassist, 2-90
- vxvg, 2-89
- vxdisk, 2-63, 2-89
- vxdiskadd, 2-89
- vxdiskconfig, 2-89
- xhost, 2-2
- xterm, 2-2
- UNIX directory subtrees, D-3
- UNIX groups
 - checking for existing oinstall group, 2-27
 - creating the dba group, 2-28
 - creating the oinstall group, 2-27
 - creating the oper group, 2-28
 - oinstall, 2-26
 - OSDBA (dba), 2-26
 - OSOPER (oper), 2-26
 - required for oracle user, 2-26
 - requirements, 2-26
 - using NIS, 2-27
- UNIX users
 - checking existence of the nobody user, 2-31
 - creating an unprivileged user on HP-UX, 2-31
 - creating the oracle user, 2-29
 - nobody, 2-26
 - oracle, 2-26
 - required for external jobs, 2-26
 - requirements, 2-26
 - setting shell limits for on Linux, 2-38
 - unprivileged user, 2-26
 - using NIS, 2-27
- UNIX workstation
 - installing from, 2-2
- unprivileged user
 - checking on AIX, Linux, Solaris, and Tru64 UNIX, 2-31
 - creating on HP-UX, 2-31
 - nobody user, 2-26
- unset command, 2-96
- unsetenv command, 2-96
- upgraded databases
 - configuring, 4-3
- upgrading, 1-11
- user names
 - ANONYMOUS, 5-3
 - BI, 5-3
 - CTXSYS, 5-3
 - DBSNMP, 5-3
 - DIP, 5-3
 - DMSYS, 5-3
 - EXFSYS, 5-3
 - HR, 5-3
 - IX, 5-3
 - LBACSYS, 5-3
 - MDDATA, 5-3
 - MDSYS, 5-3
 - MGMT_VIEW, 5-3
 - ODM, 5-3
 - ODM_MTR, 5-3
 - OE, 5-3
 - OLAPSYS, 5-3
 - ORDPLUGINS, 5-3
 - ORDSYS, 5-3
 - OUTLN, 5-3
 - PM, 5-3
 - SCOTT, 5-3
 - SH, 5-3
 - SI_INFORMTN_SCHEMA, 5-3
 - SYS, 5-3
 - SYSMAN, 5-3
 - SYSTEM, 5-3

- WK_TEST, 5-3
- WKPROXY, 5-3
- WKSYS, 5-3
- WMSYS, 5-3
- XDB, 5-3
- USER tablespace
 - raw device for
 - on AIX, 2-69
 - on HP-UX, 2-73
 - on Linux, 2-76
 - on Solaris, 2-85, 2-90
- useradd command, 2-30
- users
 - checking existence of the nobody user, 2-31
 - creating an unprivileged user on HP-UX, 2-31
 - creating the oracle user, 2-29
 - Oracle Software Owner user (oracle), 2-26
 - setting shell limits for UNIX users on Linux, 2-38
 - UNIX nobody user, 2-26
- USERS tablespace, 5-9
- UTF8 character set, 1-11
- UTLRP.SQL
 - recompiling invalid SQL modules, 4-3

V

- /var/opt/oracle/oraInst.loc file
 - reviewing log files, F-3
- Veritas Volume Manager, 2-63, 2-89
- very large databases
 - OFA naming mount points, D-4
- vfs
 - kernel subsystem on Tru64 UNIX, 2-40
- vgchange command, 2-72
- vgcreate command, 2-72
- vm
 - kernel subsystem on Tru64 UNIX, 2-40
- voldisk command, 2-66
- volume group
 - creating a directory for on HP-UX, 2-72
 - creating on AIX, 2-67
 - creating on HP-UX, 2-71
 - creating on Solaris, 2-88
- volume groups
 - creating on AIX, 2-68
 - creating on HP-UX, 2-72
 - creating on Solaris, 2-89
- volumes
 - creating logical volumes on AIX, 2-69
 - creating logical volumes on HP-UX, 2-73
 - creating logical volumes on Solaris, 2-90
- vps_ceiling parameter
 - recommended value on HP-UX, 2-35
- vxassist command, 2-90
- vxdg command, 2-89
- vxdisk command, 2-63, 2-89
- vxdiskadd command, 2-89
- vxdiskconfig command, 2-89
- VxVM, 2-63, 2-89
 - creating raw logical volumes on Solaris, 2-89

W

- WebSphere MQ
 - checking requirements on Tru64 UNIX, 2-24
 - CSD download location
 - for AIX, 2-8
 - for HP-UX, 2-12
 - for Linux, 2-18
 - for Solaris, 2-22
 - CSDs required
 - on AIX, 2-8
 - on HP-UX, 2-12
 - on Linux, 2-17
 - on Solaris, 2-22
 - on Tru64 UNIX, 2-25
 - requirement on AIX, 2-6
 - requirement on HP-UX, 2-9
 - requirement on Linux, 2-17
 - requirement on Solaris, 2-19
 - requirement on Tru64 UNIX, 2-23
- WebSphere MQ class, 4-9
- Windows
 - analogy for Oracle base directory, 2-42

X

- X emulator
 - installing from, 2-2
- X library symbolic links
 - required on HP-UX, 2-13
- X Window and X/Motif Software development
 - Motif requirement, 2-23
- X window system
 - enabling remote hosts, 2-2
- X windows
 - display errors, F-2
- X Windows Motif
 - requirement on Solaris, 2-19
- XFree86-devel
 - requirement on Linux, 2-17
- xhost command, 2-2
- xterm command, 2-2