

Oracle® Database
Installation Guide
11g Release 2 (11.2) for Linux
E10840-01

August 2009

Oracle Database Installation Guide, 11g Release 2 (11.2) for Linux

E10840-01

Copyright © 1996, 2009, Oracle and/or its affiliates. All rights reserved.

Primary Authors: Prakash Jashnani, Namrata Bhakthavatsalam, Reema Khosla

Contributing Authors: Douglas Williams, Kevin Flood, Clara Jaeckel, Emily Murphy, Terri Winters

Contributors: David Austin, Subhranshu Banerjee, Janelle Simmons, Mark Bauer, Robert Chang, Jonathan Creighton, Sudip Datta, Thirumaleshwara Hasandka, Joel Kallman, George Kotsovolos, Simon Law, Richard Long, Rolly Lv, Padmanabhan Manavazhi, Sreejith Minnanghat, Krishna Mohan, Rajendra Pingte, Hanlin Qian, Roy Swonger, Ranjith Kundapur, Aneesh Khandelwal, Barb Lundhild, Barbara Glover, Binoy Sukumaran, Hema Ramamurthy, Prasad Bagal, Martin Widjaja, Ajesh Viswambharan, Eric Belden, Sivakumar Yarlagadda, Rudregowda Mallegowda, Matthew McKerley, Trivikrama Samudrala, Akshay Shah, Sue Lee, Sangeeth Kumar, James Spiller, Saar Maoz, Rich Long, Mark Fuller, Sunil Ravindrachar, Sergiusz Wolicki, Eugene Karichkin, Joseph Francis, Srinivas Poovala, David Schreiner, Neha Avasthy, Dipak Saggi, Sudheendra Sampath, Mohammed Shah Nawaz Quadri, Shachi Sanklecha, Zakia Zerhouni, Jai Krishnani, Darcy Christensen.

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

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

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. 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, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

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

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xi
Audience	xi
Documentation Accessibility	xi
Command Syntax	xii
Accessing Documentation	xii
Related Documentation	xiii
Typographic Conventions	xiv
What's New in Oracle Database 11g	xv
New Grid Infrastructure Installation Option	xv
New Desktop and Server Class Options	xv
Daylight Savings Time Upgrade of Timestamp with Timezone Data Type	xvi
SYSASM Privilege	xvi
Fixup Scripts and Prerequisite Checks	xvi
Deinstallation Tool	xvii
Intelligent Data Placement	xvii
Oracle Automatic Storage Management Cluster File System (Oracle ACFS)	xvii
Data Pump Export and Data Pump Import	xvii
Automatic Restart of Single Instance Databases	xviii
SRVCTL Support for Single Instance Database in a Cluster	xviii
Deprecated in Oracle Database 11g Release 2 (11.2)	xviii
1 Overview of Oracle Database Installation	
New Oracle Products and Features Installed with This Release	1-1
Planning the Installation	1-1
Installing the Linux Operating System	1-2
Completing a Minimal Linux Installation	1-3
Completing a Default Linux Installation	1-3
About the Oracle Validated Configuration RPM	1-3
Installing the Oracle Validated Configuration RPM	1-4
Installation Considerations	1-4
Hardware and Software Certification	1-5
Third-Party Database Certification for SQL Developer	1-5
Multiple Oracle Homes Support	1-5
Installing the Software on a System with an Existing Oracle Installation	1-5

Grid Infrastructure for a Standalone Server.....	1-6
Oracle Cluster Synchronization Services.....	1-6
Installing Oracle Database Vault in an Oracle Data Guard Environment.....	1-6
Oracle Database Vault Default Audit Policy and Initialization Parameters	1-7
Oracle Database Installation Methods	1-7
Interactive Installation Types	1-7
Automated Installation Methods Using Response Files	1-8
Oracle Database Editions	1-8
Database Configuration Options	1-9
Preconfigured Database Types	1-9
Installation Choices that Affect Database Creation.....	1-9
Creating a Database After Installation.....	1-10
Database Storage Options	1-10
File System	1-10
Automatic Storage Management	1-11
Database Management Options	1-13
Management Options for Preconfigured Databases.....	1-14
Management Options for Custom Databases	1-14
Features Provided by Oracle Enterprise Manager Database Control	1-14
Database Backup and Recovery Options	1-15
Enabling Automated Backups.....	1-15
Backup Job Default Settings	1-16
E-mail Notification Options	1-16
Migration Consideration	1-16
Upgrade Considerations	1-16
Upgrading an Oracle Database Installed on Red Hat Enterprise Linux 2.1.....	1-16
Automatic Storage Management Is Installed with Oracle Grid Infrastructure	1-17
Daylight Savings Time Upgrade.....	1-17

2 Oracle Database Preinstallation Requirements

Logging In to the System as root	2-2
Checking the Hardware Requirements	2-3
Memory Requirements.....	2-3
System Architecture.....	2-4
Disk Space Requirements.....	2-5
Display Requirements	2-6
Run Level Requirement.....	2-6
Recommended Hardware Requirement for SQL Developer.....	2-6
Checking the Software Requirements	2-6
Operating System Requirements	2-7
Kernel Requirements	2-7
Package Requirements	2-8
Compiler Requirements	2-13
Additional Software Requirements.....	2-13
Oracle ODBC Drivers	2-14
Oracle JDBC/OCI Drivers	2-14
Oracle Messaging Gateway	2-15

Browser Requirements	2-15
Preinstallation Requirement for Oracle Database Vault	2-15
Installation Fixup Scripts.....	2-16
Checking the Network Setup.....	2-16
Installing on DHCP Computers	2-16
Installing on Multihomed Computers	2-17
Installing on Computers with Multiple Aliases	2-17
Installing on Non-Networked Computers	2-17
Creating Required Operating System Groups and Users	2-18
Creating Custom Configuration Groups and Users for Job Roles.....	2-18
Users for Oracle Installations with Job Role Separation	2-18
Database Groups for Job Role Installations	2-19
Grid Infrastructure Groups for Job Role Installations.....	2-19
Creating Database Operating System Groups and Users with Job Role Separation	2-20
Creating the Oracle Inventory Group.....	2-20
Creating the OSDBA Group for Database Installations	2-21
Creating an OSOPER Group for Database Installation.....	2-22
Creating the OSASM Group for Automatic Storage Management	2-22
Creating the OSDBA Group for Oracle Grid Infrastructure.....	2-22
Creating the OSOPER Group for Automatic Storage Management	2-22
Creating the Oracle Software Owner User.....	2-22
Determining if an Oracle Software Owner User Exists.....	2-22
Creating an Oracle Software Owner User.....	2-23
Modifying an Existing Oracle Software Owner User.....	2-23
Configure Oracle Installation Owner Shell Limits	2-23
Configuring Kernel Parameters.....	2-25
Displaying and Changing Kernel Parameter Values	2-26
Identifying Required Software Directories	2-28
Oracle Base Directory	2-28
Oracle Inventory Directory.....	2-29
Oracle Home Directory	2-30
Identifying or Creating an Oracle Base Directory	2-30
Identifying an Existing Oracle Base Directory.....	2-30
Creating an Oracle Base Directory.....	2-31
Choosing a Storage Option for Oracle Database and Recovery Files.....	2-32
Creating Directories for Oracle Database or Recovery Files	2-33
Guidelines for Placing Oracle Database Files on a File System.....	2-33
Creating Required Directories.....	2-34
Configuring Storage for Oracle Database Files Using Block Devices.....	2-35
Configuring Disk Devices for Oracle Database.....	2-35
Example of Creating a Udev Permissions File for Oracle Database.....	2-36
Example of Configuring Block Device Storage for Oracle Database.....	2-36
Stopping Existing Oracle Processes.....	2-37
Configuring the oracle User's Environment	2-38

3 Oracle Grid Infrastructure

Requirements for Oracle Grid Infrastructure Installation.....	3-2
---	-----

Memory Requirements.....	3-2
Disk Space Requirements.....	3-3
Configuring the User's Environment.....	3-3
Oracle ACFS and Oracle ADVM Support.....	3-3
Migrating Existing Automatic Storage Management Instances.....	3-4
Automatic Storage Management Installation Considerations.....	3-5
Preparing Disks for an Automatic Storage Management Installation.....	3-5
General Steps for Configuring Automatic Storage Management.....	3-6
Step 1: Identifying Storage Requirements for Automatic Storage Management.....	3-6
Step 2: Using an Existing Automatic Storage Management Disk Group.....	3-9
Step 3: Creating DAS or SAN Disk Partitions for Automatic Storage Management.....	3-10
Step 4: Configuring Disks for Automatic Storage Management.....	3-10
Configuring Disks for Automatic Storage Management Using the Automatic Storage Management Library Driver (ASMLIB).....	3-11
Installing and Configuring Grid Infrastructure for a Standalone Server.....	3-14
Installing Grid Infrastructure with a New Database Installation.....	3-14
Installing Grid Infrastructure for an Existing Database.....	3-18
Manually Configuring Oracle Automatic Storage Management Disk Groups.....	3-18
Testing the Automatic Storage Management Installation.....	3-19

4 Installing Oracle Database

Preinstallation Considerations.....	4-1
Performing Multiple Oracle Database Installations in Response File or Silent Mode.....	4-1
Reviewing Component-Specific Installation Guidelines.....	4-1
Selecting the Database Character Set.....	4-2
Accessing the Installation Software.....	4-3
Downloading Oracle Software.....	4-3
Downloading the Installation Archive Files from OTN.....	4-4
Downloading the Software from Oracle E-Delivery.....	4-4
Extracting the Installation Files.....	4-5
Copying the Software to the Hard Disk.....	4-5
Mounting Disks.....	4-5
Copying the Oracle Database Software to a Hard Disk.....	4-6
Database Security Options.....	4-7
Installing the Oracle Database Software.....	4-8
Running Oracle Universal Installer.....	4-8
Installing Oracle Database Examples.....	4-15

5 Oracle Database Postinstallation Tasks

Required Postinstallation Tasks.....	5-1
Downloading and Installing Patches.....	5-1
Recommended Postinstallation Tasks.....	5-2
Creating a Backup of the root.sh Script.....	5-2
Configuring New or Upgraded Databases.....	5-3
Creating and Configuring Additional Operating System Accounts.....	5-3
Configuring the Accounts of Oracle Users.....	5-3
Setting the NLS_LANG Environment Variable.....	5-4

Generating the Client Static Library	5-4
Guidelines for Setting Semaphore Parameters	5-4
Create a Fast Recovery Area Disk Group	5-5
About the Fast Recovery Area and the Fast Recovery Area Disk Group	5-5
.....Creating the Fast Recovery Area Disk Group	5-5
Required Product-Specific Postinstallation Tasks	5-6
Configuring Oracle Net Services	5-6
Configuring Oracle Label Security	5-7
Configuring Oracle Database Vault	5-7
Configuring Oracle Messaging Gateway	5-7
Modifying the listener.ora File for External Procedures	5-8
Modifying the tnsnames.ora File for External Procedures	5-8
Setting Up the mgw.ora Initialization File	5-9
Configuring Oracle Precompilers	5-9
Configuring Pro*C/C++	5-9
Configuring Pro*FORTRAN	5-10
Configuring Secure Sockets Layer	5-10
Installing Oracle Text Supplied Knowledge Bases	5-10
Configuring or Reinstalling Oracle XML DB	5-10
Configuring and Using Direct NFS Client	5-11
Direct NFS Client	5-11
Enabling Direct NFS Client	5-11
Disabling Direct NFS Client	5-13
Checking NFS Buffer Size Parameters	5-13
Postinstallation Tasks for SQL Developer	5-13
Migrating User Settings from Release 1.0	5-13
Migrating Information from Previous Releases	5-14
Location of User-Related Information	5-14

6 Getting Started with Oracle Database

Checking the Installed Oracle Database Contents and Directory Location	6-1
Logging In to Oracle Enterprise Manager Database Control	6-1
Managing Automatic Storage Management	6-3
Starting and Stopping Automatic Storage Management	6-3
Automatic Storage Management Utilities	6-3
Accessing Oracle Database with SQL*Plus	6-4
Accessing Oracle Database with SQL Developer	6-4
Reviewing Accounts and Passwords	6-5
Unlocking and Resetting User Passwords	6-8
Using Database Control to Unlock Accounts and Reset Passwords	6-8
Using SQL*Plus to Unlock Accounts and Reset Passwords	6-9
Unlocking and Changing Passwords	6-9
Identifying Databases	6-10
Locating the Server Parameter File	6-11
Reviewing Tablespaces and Data Files, Redo Log Files, and Control Files	6-11
Identifying Tablespaces and Data Files	6-11
Locating Redo Log Files	6-12

Locating Control Files.....	6-13
-----------------------------	------

7 Removing Oracle Database Software

Reconfiguring Oracle Cluster Synchronization Services.....	7-1
Removing Oracle Software Using the Deinstallation Tool.....	7-1
About the Deinstallation Tool.....	7-2
Example of Running the Deinstall Command for an Oracle Database.....	7-3
Example of a Deinstallation Parameter File for Oracle Database.....	7-4
Example of a Deinstallation Parameter File for Oracle Grid Infrastructure for a Standalone Server 7-4	
Removing Oracle Software.....	7-5

A Installing and Configuring Oracle Database Using Response Files

How Response Files Work?.....	A-1
Reasons for Using Silent Mode or Response File Mode.....	A-2
Creating a Database Using Automatic Storage Management as the Storage Option for Database Files A-3	
General Procedure for Using Response Files.....	A-3
Creating the oraInst.loc File.....	A-3
Preparing a Response File.....	A-4
Editing a Response File Template.....	A-4
Recording a Response File.....	A-5
Running Oracle Universal Installer Using a Response File.....	A-7
Running Net Configuration Assistant Using a Response File.....	A-8
Running Database Configuration Assistant Using a Response File.....	A-8
Using Database Configuration Assistant in Response File Mode.....	A-9
Using Database Configuration Assistant in Silent Mode.....	A-9
Running Database Configuration Assistant in Response File or Silent Mode.....	A-9
Postinstallation Configuration Using a Response File.....	A-10
About the Postinstallation Configuration File.....	A-10
Running Postinstallation Configuration Using a Response File.....	A-11

B Cloning an Oracle Home

Cloning an Oracle Home.....	B-1
Configuring Oracle Configuration Manager in a Cloned Oracle Home.....	B-3

C Using NAS Devices

General Configuration Guidelines for NAS Devices.....	C-1
NFS Feature Description.....	C-1
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 Automatic Storage Management.....	C-4
NFS Mount Options.....	C-5

D Optimal Flexible Architecture

Overview of the Optimal Flexible Architecture Standard	D-1
Implementing Optimal Flexible Architecture	D-1
File Systems	D-1
Number of File Systems	D-2
Naming Conventions	D-2
Naming Directories.....	D-2
Oracle Base Directory Naming Convention.....	D-2
Naming Mount Points for Very Large Databases (VLDBs).....	D-3
Referring to Path Names.....	D-3
Oracle Home Directory Naming Convention.....	D-3
Naming Subdirectories	D-3
Naming Database Files.....	D-4
Separating Segments with Different Requirements.....	D-5
Exploiting the Optimal Flexible Architecture Structure for Oracle Files.....	D-6
Optimal Flexible Architecture File Mapping	D-6

E Managing Oracle Database Port Numbers

About Managing Ports	E-1
Viewing Port Numbers and Access URLs	E-2
Port Numbers and Protocols of Oracle Components	E-2
Changing the Oracle Enterprise Management Agent Port.....	E-3
Changing the Oracle Enterprise Manager Database Control Ports.....	E-3
Changing the Oracle XML DB Ports.....	E-4

F Configuring Oracle Database Globalization Support

Installing and Using Oracle Components in Different Languages.....	F-1
Configuring Oracle Components to Run in Different Languages	F-1
Determining the Operating System Locale by Using the LANG Environment Variable	F-2
Configuring Locale and Character Sets by Using the NLS_LANG Environment Variable	F-3
Installing Translation Resources	F-3
Running Oracle Universal Installer in Different Languages	F-4

G Troubleshooting

Verify Requirements.....	G-1
X Window Display Errors	G-2
Remote Terminal Installation Error	G-2
What to Do If an Installation Error Occurs?.....	G-3
Reviewing the Log of an Installation Session	G-3
Troubleshooting Host Name Changes and CSS.....	G-4
Troubleshooting Configuration Assistants	G-4
Configuration Assistant Failure	G-4
Irrecoverable Errors	G-5
Troubleshooting Inventory Issues	G-5

Troubleshooting Screen Display Issues.....	G-5
Silent-Mode Response File Error Handling.....	G-5
Cleaning Up After a Failed Installation.....	G-6

H Frequently Asked Questions About Installation

Installing Oracle Database or Oracle Database Client	H-1
Installing Oracle Database Tools	H-3
Installing Oracle Database with Oracle Applications	H-8
Installing Oracle Database Heterogeneous Connectivity Tools (Gateways).....	H-9

I Country Codes

Valid Country Codes.....	I-1
--------------------------	-----

Glossary

Index

Preface

This guide provides instructions about installing and configuring Oracle Database for Linux. This guide covers Optimal Flexible Architecture, Database Storage Options, and Database Configuration Options. This guide also talks about installing and configuring database using response files, globalization support, ports, and troubleshooting.

The preface contains the following topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Command Syntax](#)
- [Accessing Documentation](#)
- [Related Documentation](#)
- [Typographic Conventions](#)

Audience

Oracle Database Installation Guide for Linux guide is intended for anyone responsible for installing Oracle Database 11g Release 2 (11.2) on a single Linux x86 system. Additional installation guides for Oracle Database, Oracle Real Application Clusters, Oracle Clusterware, Oracle Database Examples, and Oracle Enterprise Manager Grid Control are available on the relevant installation media.

See Also: *Oracle Database Installation Guide for Linux* to install Oracle Database using the default settings

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. 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. Accessibility 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 more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers 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, some screen readers 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.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at <http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

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/rdskc0t1d0s6 of=/dev/rst0 bs=10b \ count=10000</pre>
braces { }	Braces indicate required items: <pre>.DEFINE {macro}</pre>
brackets []	Brackets indicate optional items: <pre>cvtcrt <i>termname</i> [<i>outfile</i>]</pre>
ellipses ...	Ellipses indicate an arbitrary number of similar items: <pre>CHKVAL <i>fieldname</i> <i>value1</i> <i>value2</i> ... <i>valueN</i></pre>
<i>italics</i>	Italic type indicates a variable. Substitute a value for the variable: <pre><i>library_name</i></pre>
vertical line	A vertical line indicates a choice within braces or brackets: <pre>FILE <i>filesize</i> [K M]</pre>

Accessing Documentation

The documentation for this release 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.

This guide contains information required to install Oracle Database 11g Release 2 (11.2) on various platforms of Linux. Ensure that you review information related to the platform on which you intend to install Oracle Database 11g.

The platform-specific documentation for this product is available in both Adobe portable document format (PDF) and HTML format on the product media. To access the platform-specific documentation on media:

1. Use a Web browser to open the `welcome.htm` file in the top-level directory of the media.
2. For DVD 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 11g products is available in both HTML and PDF formats in the Oracle Database 11g Release 2 (11.2) Online Documentation Library. To check for updates to this document and to view other Oracle documentation, select the Documentation link or the Software & Patches link on the Oracle Database 11g Release 2 Web site.

Related Documentation

The platform-specific documentation for Oracle Database 11g products includes the following manuals:

- *Oracle Database Concepts*
- *Oracle Database New Features Guide*
- *Oracle Database Licensing Information*
- *Oracle Database Readme*
- *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX*
- *Oracle Grid Infrastructure Installation Guide*
- *Oracle Database Release Notes for Linux*
- *Oracle Database Client Installation Guide for Linux*
- *Oracle Database Examples Installation Guide*
- *Oracle Real Application Clusters Installation Guide*
- *Oracle Database Quick Installation Guide for Linux x86-64*
- *Oracle Database Client Quick Installation Guide for Linux x86-64*
- *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems*
- *Oracle Database Storage Administrator's Guide*
- *Oracle Database Upgrade Guide*

- *Oracle Database 2 Day DBA*
- *Oracle Application Express Installation Guide*

For information about Oracle error messages, see *Oracle Database Error Messages*. Oracle error message documentation is available only in HTML. If you only have access to the Oracle Database 11g Release 2 (11.2) Online Documentation Library, then you can browse the error messages by range. Once you find the specific range, use your browser's "find in page" feature to locate the specific message. When connected to the Internet, you can search for a specific error message using the error message search feature of the Oracle online documentation.

Many books in the documentation set use the sample schemas of the seed database, which is installed by default when you install Oracle. Refer to *Oracle Database Sample Schemas* for information on how these schemas were created and how you can use them yourself.

Printed documentation is available for sale in the Oracle Store at:

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

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

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

If you already have a user name and password for Oracle Technology Network, then you can go directly to the documentation section of the Oracle Technology Network Web site at:

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

Refer to *Oracle Database Release Notes for Linux* or important information that was not available when this book was released. The release notes for Oracle Database 11g are updated regularly. The most recent version is available on Oracle Technology Network at:

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

Typographic Conventions

The following text conventions are used in this document:

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

What's New in Oracle Database 11g

The following is a list of new features or enhancements provided with Oracle Database 11g Release 2 (11.2):

- [New Grid Infrastructure Installation Option](#)
- [SYSASM Privilege](#)
- [Fixup Scripts and Prerequisite Checks](#)
- [New Desktop and Server Class Options](#)
- [Daylight Savings Time Upgrade of Timestamp with Timezone Data Type](#)
- [Deinstallation Tool](#)
- [Intelligent Data Placement](#)
- [Oracle Automatic Storage Management Cluster File System \(Oracle ACFS\)](#)
- [Data Pump Export and Data Pump Import](#)
- [Automatic Restart of Single Instance Databases](#)
- [SRVCTL Support for Single Instance Database in a Cluster](#)
- [Deprecated in Oracle Database 11g Release 2 \(11.2\)](#)

New Grid Infrastructure Installation Option

Oracle Database 11g Release 2 introduces the Grid Infrastructure installation. For single instance databases, the grid infrastructure includes Automatic Storage Management, the listener, and Oracle Restart. Oracle Restart is a new feature that provides the ability to monitor, manage, and automatically restart on failure the Oracle Database environment including the Oracle Database instance, Oracle Automatic Storage Management instance, and listeners. In a clustered environment, the Grid Infrastructure includes Oracle Clusterware, Automatic Storage Management, and the listener.

If you want to use grid infrastructure for a standalone server, then you must install the oracle software from the grid infrastructure media before you install the database.

See Also: ["Oracle Grid Infrastructure"](#)

New Desktop and Server Class Options

Oracle Database 11g Release 2 introduces a new option that allows you to specify the type of system on which the database is installed. If you are installing on a laptop or a

desktop, then select the Desktop Class option; otherwise, select the Server Class option to install on a server. These options are available on the System Class screen.

There is no difference in the software that gets installed after you select any one option but the Desktop Class option installs a single instance database without the advanced configuration options.

See Also: ["Interactive Installation Types"](#) for more information about the desktop and the server class options

Daylight Savings Time Upgrade of Timestamp with Timezone Data Type

When time zone version files are updated due to Daylight Saving Time changes, `TIMESTAMP WITH TIMEZONE (TSTZ)` data could become stale. In previous releases, database administrators ran the SQL script `ut1tzuv2.sql` to detect TSTZ data affected by the time zone version changes and then had to carry out extensive manual procedures to update the TSTZ data.

With this release, TSTZ data is updated transparently with very minimal manual procedures using newly provided `DBMS_DST` PL/SQL packages. In addition, there is no longer a need for clients to patch their time zone data files.

See Also:

- *Oracle Database Upgrade Guide* for information about preparing to upgrade Timestamp with Time Zone data
- *Oracle Database Globalization Support Guide* for information about how to upgrade the Time Zone file and Timestamp with Time Zone data
- *Oracle Call Interface Programmer's Guide* for information about performance effects of clients and servers operating with different versions of Time Zone files

SYSASM Privilege

Starting with 11g release 2, Oracle ASM administration must be done with the `SYSASM` privilege. The `SYSASM` privilege also can be granted using password authentication on the Oracle ASM instance.

You can designate `OPERATOR` privileges (a subset of the `SYSASM` privileges, including starting and stopping Oracle ASM) to members of the `OSOPER` for Oracle ASM group.

Using the `SYSASM` privilege for ASM administration creates a clearer division of responsibility between ASM administration and database administration. It also provides the optional capability to prevent different databases using the same storage from accidentally overwriting each other's files.

See Also: *Oracle Database Storage Administrator's Guide* for more information about `SYSASM` privilege and `OSASM` operating system group

Fixup Scripts and Prerequisite Checks

With this release, the installer (OUI) detects when minimum requirements for installation are not completed, and creates scripts, called fixup scripts, to resolve many

incomplete system configuration requirements. If OUI detects an incomplete task, then generate the fixup script by clicking the Fix & Check Again button.

For Oracle Clusterware, you also can have Cluster Verification Utility (CVU) generate fixup scripts before installation.

The fixup script is generated during installation. You are prompted to run the script as root in a separate terminal session. When you run the script, it sets some of the system parameters to Oracle-recommended values, if necessary, and completes other operating system configuration tasks.

Deinstallation Tool

Oracle Universal Installer no longer removes Oracle software. Use the new Deinstallation Tool (`deinstall`) available on the installation media before installation as well as in the Oracle home directory after installation.

Run this tool from `ORACLE_HOME` first. It is located in:

```
$ORACLE_HOME/deinstall.
```

An error occurs if there are any missing files and the tool prompts you to download the standalone version available in the `deinstall` directory on the installation media.

See Also: ["Removing Oracle Software"](#)

Intelligent Data Placement

The Intelligent Data Placement feature enables you to specify disk regions on Oracle ASM disks to ensure that frequently accessed data is placed on the hot regions (outermost tracks) which provide higher performance.

See Also: *Oracle Database Storage Administrator's Guide* for more information about Automatic Storage Management Intelligent Data Placement

Oracle Automatic Storage Management Cluster File System (Oracle ACFS)

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) is a new multi-platform, scalable file system, and storage management design that extends Oracle ASM technology to support data which cannot be stored in Oracle ASM, in both single instance and cluster configurations. Additionally, Oracle ACFS provides snapshot functionality for a point in time copy of an Oracle ACFS system.

The software required for Oracle Automatic Storage Management Cluster File System is installed with the grid infrastructure installation.

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

Data Pump Export and Data Pump Import

Data Pump provides a legacy mode in which you can use original Export and Import parameters when performing Data Pump Export and Import operations.

See Also: *Oracle Database Utilities* for more information about Data Pump Legacy Mode

Automatic Restart of Single Instance Databases

Oracle Restart is a new feature included in this release to enhance the availability of Oracle databases in a single-instance environment. If you install Oracle Restart, then in the event of a temporary failure of any part of the Oracle Database software stack, including the database, listener, and Automatic Storage Management instance, Oracle Restart automatically restarts the failed component. In addition, Oracle Restart starts all these components when the database host computer is restarted. The components are started in the proper order, taking into consideration the dependencies among components.

See Also: Chapter 4, "Configuring Automatic Restart of an Oracle Database" in the *Oracle Database Administrator's Guide* for more information about Oracle Restart

New Method of Installing Oracle Automatic Storage Management

In past releases, Oracle Automatic Storage Management was installed as part of the Oracle Database installation. With Oracle Database 11g Release 2 (11.2), Oracle Automatic Storage Management is installed when you install the grid infrastructure components and shares an Oracle home with Oracle Clusterware when installed in a cluster such as with Oracle RAC or with Oracle Restart on a single instance database.

If you want to upgrade an existing Oracle ASM, then you must upgrade Oracle ASM by running an Oracle grid infrastructure upgrade. If you do not have Oracle ASM installed and you want to use Oracle ASM as your storage option, then you must complete an Oracle grid infrastructure installation before you start your Oracle Database installation.

See Also: "Oracle grid infrastructure" for more information about installing the Oracle grid infrastructure software

SRVCTL Support for Single Instance Database in a Cluster

SRVCTL has been enhanced to support single instance databases with Oracle Restart on standalone servers and on clusters with Oracle Clusterware. SRVCTL is a command-line interface used to manage Oracle processes (database instance, listener, Oracle ASM instance) when using Oracle Restart. With SRVCTL, you can manage the Oracle Restart configuration, see the status of processes managed by Oracle Restart, and start or stop processes such as the Oracle Database.

See Also: Chapter 4, "Configuring Automatic Restart of an Oracle Database" in the *Oracle Database Administrator's Guide* for more information about SRVCTL commands

Deprecated in Oracle Database 11g Release 2 (11.2)

The following are not supported or not available anymore with Oracle Database 11g Release 2:

- Installing data files directly on raw devices is no longer available during installation with Oracle Universal Installer or Database Configuration Assistant. You must use a file system, or use Automatic Storage Management.
- Oracle Ultra Search

Overview of Oracle Database Installation

This chapter describes the different installation types of Oracle Database for Linux and issues to consider before you install Oracle Database:

- [New Oracle Products and Features Installed with This Release](#)
- [Planning the Installation](#)
- [Installing the Linux Operating System](#)
- [Installation Considerations](#)
- [Oracle Database Installation Methods](#)
- [Oracle Database Editions](#)
- [Database Configuration Options](#)
- [Database Storage Options](#)
- [Database Management Options](#)
- [Database Backup and Recovery Options](#)
- [E-mail Notification Options](#)
- [Migration Consideration](#)
- [Upgrade Considerations](#)

New Oracle Products and Features Installed with This Release

There are many new features and products installed with this release. Refer to [What's New in Oracle Database 11g](#) for more information on the new features and products installed with this release.

Planning the Installation

The Oracle Database installation process consists of the following phases:

1. **Read the release notes:** Read *Oracle Database Release Notes for Linux* before you begin the installation. The release notes are available with the platform-specific documentation. The latest version of the release notes is available on Oracle Technology Network at:
<http://www.oracle.com/technology/documentation>
2. **Review the licensing information:** Although the installation media in the media pack contain many Oracle components, you are permitted to use only those components for which you have purchased licenses.

Oracle Support Services does not provide support for components for which licenses have not been purchased.

See Also: *Oracle Database Licensing Information* for more details

3. **Plan the installation:** This chapter describes the Oracle products that you can install and issues that you must consider before starting the installation.
You can also refer to [Appendix H](#), which covers frequently asked questions about installing Oracle Database components, such as how to install Oracle Database if the site uses Oracle applications or if you need multiple Oracle Database connections.
4. **Complete preinstallation tasks:** [Chapter 2](#) describes preinstallation tasks that you must complete before installing the product.
5. **Install the software:** Use the following sections to install Oracle Database and Oracle grid infrastructure for a standalone server:
 - [Chapter 4](#) describes how to use Oracle Universal Installer to install Oracle Database. Also describes how to clone an Oracle home.
 - [Chapter 3](#) describes how to install Oracle grid infrastructure for a standalone server.
 - [Chapter 7](#) describes how to remove Oracle Database software.
 - [Appendix A](#) provides information on performing silent, or response file installations, which you may want to use if you must perform multiple installations of Oracle Database.
 - [Appendix B](#) provides information on cloning Oracle home.
 - [Appendix F](#) describes globalization support information.
 - [Appendix G](#) provides troubleshooting advice in case you encounter problems with the installation.
6. **Complete postinstallation tasks:** [Chapter 5](#) describes recommended and required postinstallation tasks.
7. **Get started using Oracle Database:** Use the following sections to get started with Oracle Database:
 - [Chapter 6](#) describes how to check the contents of the installed Oracle Database, how to start various tools, and how to locate various files.
 - [Appendix C](#) describes the network attached storage devices, which you can use to store Oracle database files and Oracle software.
 - [Appendix E](#) explains the method to manage Oracle Database port numbers.

Installing the Linux Operating System

This section provides information about installing a supported Linux distribution. It contains the following topics:

- [Completing a Minimal Linux Installation](#)
- [Completing a Default Linux Installation](#)
- [About the Oracle Validated Configuration RPM](#)
- [Installing the Oracle Validated Configuration RPM](#)

Completing a Minimal Linux Installation

To complete a minimal Linux installation, select a minimal install option (either a custom installation where you select the Minimal option from Package Group Selection, or where you deselect all packages except for the Base pack). This installation lacks many RPMs required for installation. However, when you install the Oracle Validated RPM for your platform, the RPM downloads the minimum number of packages required to run Oracle Clusterware and Oracle Database.

Unbreakable Linux Network (ULN) customers can obtain the Oracle Validated RPM by using `up2date`. If you are not a ULN customer, and your operating system is Red Hat or Oracle Enterprise Linux, then you can obtain the Oracle Validated RPM at the following URLs:

Enterprise Linux 4:

<http://oss.oracle.com/el4/oracle-validated/>

Enterprise Linux 5:

<http://oss.oracle.com/el5/oracle-validated/>

Note: The Oracle Validated RPM installs the X11 client libraries, but does not install the X Window System server packages. To use graphical user interfaces such as Oracle Universal Installer, configuration assistants, and Enterprise Manager, set the display to a system with X Window System server packages.

Completing a Default Linux Installation

If you do not install the Oracle Validated RPM, then Oracle recommends that you install your Linux operating system with the default software packages (RPMs). This installation includes most of the required packages and helps you limit manual checks of package dependencies. Oracle recommends that you do not customize the RPMs during installation.

For information about a default installation, log on to My Oracle Support (formerly *OracleMetaLink*):

<https://metalink.oracle.com>

Search for "Default RPM."

After installation, start `system-config-packages` and add all the Legacy Software Development packages. Review system requirements for the distribution to ensure that you have all required kernel packages installed, and complete all other configuration tasks required for the distribution, and for the system configuration.

About the Oracle Validated Configuration RPM

If the Linux distribution is Oracle Enterprise Linux, or Red Hat Enterprise Linux, and you are an Unbreakable Linux customer, then you can complete most preinstallation configuration tasks by using the Oracle Validated Configurations Setup packages, available from the Unbreakable Linux Network (ULN).

When Unbreakable Linux Network is installed, the Oracle Validated Configuration packages sets and verifies system parameters based on recommendations from the Oracle Validated Configurations program, and installs any additional packages needed for installing Oracle Clusterware and Oracle Database. It also updates

`sysctl.conf` settings, system startup parameters, user limits, and driver parameters to values which extensive testing shows will provide better performance.

To become an Oracle Unbreakable Linux customer, contact your sales representative, or purchase a license from the Unbreakable Linux store:

http://oraclestore.oracle.com/OA_HTML/ibeCZzpHome.jsp

To register your server on the Unbreakable Linux Network, or to find out more information, refer to the following URL:

<https://linux.oracle.com>

Installing the Oracle Validated Configuration RPM

Use the following procedure to subscribe to Oracle Unbreakable Linux channels, and to add the Oracle Software for Enterprise Linux channel that distributes the Oracle Validated Configurations Setup RPM:

1. Complete a default Oracle Enterprise Linux workstation installation, or a default Red Hat Enterprise Linux installation.
2. Register the server with Unbreakable Linux Network (ULN). By default, you are registered for the Enterprise Linux Latest channel for the operating system and hardware.
3. Log in to ULN at the following URL:
<https://linux.oracle.com>
4. Click the **Systems** tab, and in the System Profiles list, select a registered server. The System Details window opens, and displays the subscriptions for the server.
5. From the Available Channels list, select the **Oracle Software for Enterprise Linux channel** that is appropriate for the installation of Linux. For example: Oracle Software for Enterprise Linux 4 (x86-64).
6. Click **Subscribe**.
7. From a terminal session, as a `root` user, enter the following command:

```
# up2date --nox --show-channels
```

An output similar to the following is displayed. This output indicates that you have subscribed to the Oracle Software for Enterprise Linux channel. For example:

```
e14_i386_latest  
e14_i386_oracle
```

8. Open a terminal session as a `root` user, and install the Oracle Validated Configurations Setup RPM with `up2date` with the following command:

```
# up2date --install oracle-validated
```

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 Support](#)

- [Grid Infrastructure for a Standalone Server](#)
- [Oracle Cluster Synchronization Services](#)
- [Installing Oracle Database Vault in an Oracle Data Guard Environment](#)
- [Oracle Database Vault Default Audit Policy and Initialization Parameters](#)

Hardware and Software Certification

The platform-specific hardware and software requirements included in this guide were current when this guide was published. However, because new platforms and operating system software versions might be certified after this guide is published, review the certification matrix on the My Oracle Support (formerly Oracle*MetaLink*) Web site for the most up-to-date list of certified hardware platforms and operating system versions. The My Oracle Support (formerly Oracle*MetaLink*) Web site is available at the following URL:

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

You must register online before using My Oracle Support (formerly Oracle*MetaLink*). After logging in, click **Certify** on the top right-hand side of the screen. The Certifications page appears. You can choose the certification either by product or by platform.

Other options include Oracle's Certification Matrices, Desupport Notices, and Product Availability.

Third-Party Database Certification for SQL Developer

SQL Developer can be used to view metadata and data of several non-Oracle databases. The following table lists the third-party database certifications.

Database	Releases	Notes
Microsoft Access	Access 97 Access 2000 Access 2003	For any Access release: no JDBC driver needed, but you must ensure read access to the system tables in the .mdb file.
Microsoft SQL Server	SQL Server 7 SQL Server 2000 SQL Server 2005	For any Microsoft SQL Server release: JDBC driver <code>jtods-1.2.2.jar</code> required. This is included in the <code>jtods-1.2-dist.zip</code> file available from sourceforge.net
MySQL	MySQL 3.x MySQL 4.x MySQL 5.x	For any MySQL release: JDBC driver required. For MySQL 5.x: <code>mysql-connector-java-5.0.4-bin.jar</code> is required, which is included in <code>mysql-connector-java-5.0.4.zip</code>

Multiple Oracle Homes Support

This product supports multiple Oracle homes. This means, you can install this release or earlier 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 Oracle Database 11g Release 2 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, then the installation fails.

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

Grid Infrastructure for a Standalone Server

The grid infrastructure for a standalone server provides the infrastructure to include your single instance database in an enterprise grid architecture. Oracle Database 11g Release 2 (11.2) combines these infrastructure products into one software installation called the grid infrastructure home. On a single instance database, the grid infrastructure home includes Oracle Restart and Automatic Storage Management (ASM) software.

If you want to use Automatic Storage Management or Oracle Restart, then you should first install Oracle grid infrastructure for a standalone server, and then install Oracle Database 11g Release 2 (11.2).

See Also: [Chapter 3, "Oracle Grid Infrastructure"](#) for more information about installing Oracle grid infrastructure for a standalone server

Oracle Cluster Synchronization Services

When you install Oracle grid infrastructure for a standalone server, it will configure the single-node version of Oracle Cluster Synchronization Services (CSS). The CSS service is required to enable synchronization between an Automatic Storage Management instance and the database instances that rely on it for database file storage. Because the service must be running before any Automatic Storage Management instance or database instance starts, it is configured to start automatically by Oracle Restart before the Automatic Storage Management instance is started.

For Oracle RAC installations, the CSS daemon is installed with Oracle Clusterware in a separate Oracle home directory (also called the grid infrastructure directory). For single instance installations, the CSS daemon is installed in and runs from the grid infrastructure home which is the same home that runs Automatic Storage Management.

See Also:

- ["Automatic Storage Management"](#) on page 1-11
- ["Reconfiguring Oracle Cluster Synchronization Services"](#) on page 7-1
- ["Removing Oracle Software"](#)

Installing Oracle Database Vault in an Oracle Data Guard Environment

If you plan to use Oracle Data Guard with Oracle Database Vault, then refer to Note 754065.1 on the My Oracle Support (formerly *OracleMetaLink*) Web site at the following URL:

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

Oracle Database Vault Default Audit Policy and Initialization Parameters

Oracle Database Vault installs a baseline database auditing policy. This policy covers the access control configuration information stored in Database Vault database tables, information stored in Oracle Catalog (rollback segments, tablespaces, and so on), the use of system privileges, and Oracle Label Security configuration. When you install Oracle Database Vault, the security specific database initialization parameters are initialized with default values.

See Also: *Oracle Database Vault Administrator's Guide* for more information on the database audit policy

Oracle Database Installation Methods

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

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

Interactive Installation Types

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

Starting with Oracle Database 11g Release 2 (11.2), Oracle Universal Installer provides you the following options:

- **Desktop Class:** Select this option if you are installing on a laptop or desktop class system. This option includes a starter database and allows minimal configuration. This option is designed for those who want to quickly set up a database.
- **Server Class:** Select this option if you are installing on a server class system, such as what you would use when deploying Oracle in a production data center. This option allows for more advanced configuration options. Advanced configuration options available with this option include Oracle RAC, Automatic Storage Management, backup and recovery configuration, integration with Enterprise Manager Grid Control, and more fine-grained memory tuning, among others.

Furthermore, the Server Class option provides you with the following installation types:

- **Typical:** Select this installation method if you want to quickly install Oracle Database. This installation type requires minimal user input. It installs the software and optionally creates a general-purpose database using the information that you specify on the screen. It is the default installation type.
- **Advanced:** Select this installation type if you want to complete any of the following tasks:
 - 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
 - Specify different passwords for administrative schemas
 - Configure automated backups or Oracle Enterprise Manager notifications
 - Configure Oracle Configuration Manager

- Customize components from the available components list. In the Select Database Edition screen, if you select **Enterprise Edition**, then Oracle Universal Installer automatically selects the components most customers need for their Oracle Database installation. You can also click **Select Options** to customize components from the components list.

See Also: ["Reviewing Component-Specific Installation Guidelines"](#) on page 4-1 for additional information on Oracle database installation

Automated Installation Methods Using Response Files

By creating a response file and specifying this file when you start Oracle Universal Installer, you can automate some or all of the Oracle Database installation. These automated installation methods are useful if you must 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 Oracle Universal Installer in the following modes, depending on whether you specify all of the required information or not:

- **Silent Mode:** Oracle Universal Installer runs in silent mode if you use a response file that specifies all required information, and specify the `-silent` option when starting Oracle Universal Installer. None of the Oracle Universal Installer screens are displayed.
- **Response File Mode:** Oracle Universal Installer runs in response file mode if you do not specify all required information in the response file.

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

Oracle Database Editions

You can choose one of the following database editions when installing Oracle Database 11g Release 2 (11.2):

- **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. This option also allows you to enable or disable individual components from a components list.
- **Standard Edition:** This installation type is designed for department or workgroup-level applications and for small and medium-sized enterprises (SMEs). It is engineered to provide core relational database management services and options. It installs an integrated set of management tools, full distribution, replication, Web features, and facilities for building business-critical applications.
- **Standard Edition One:** This installation type is designed for department, workgroup-level, or web applications. From single instance environments for small business to highly distributed branch environments, Oracle Database Standard Edition One includes all the facilities necessary to build business-critical applications.

See Also:

- You must install Oracle Database Client separately. You cannot install it during an Oracle Database installation. Refer to *Oracle Database Client Installation Guide for Linux* for installation instructions.
- *Oracle Database Licensing Information* for more information about the features available with each Oracle Database edition and for information about licensing

Note:

- The installation process is the same for all the database editions.
 - Ensure that you install only those products for which you have a valid license.
-
-

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, then Oracle Universal Installer uses Oracle Database Configuration Assistant 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 the requirements.

This section describes the following database configuration options:

- [Preconfigured Database Types](#)
- [Installation Choices that Affect Database Creation](#)
- [Creating a Database After Installation](#)

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

Refer to the online help provided by either Oracle Universal Installer or Oracle Database Configuration Assistant for a description of these preconfigured database types.

Installation Choices that Affect Database Creation

Oracle Universal Installer runs Oracle Database Configuration Assistant in one of two modes, depending on the choices that you make during the installation:

- Silent or response file mode

If you choose either the Enterprise Edition or Standard Edition database edition, then choose to create a preconfigured database type. Oracle Universal Installer prompts you for the minimum amount of information required to create a database of the type you choose. It then runs Oracle Database Configuration

Assistant in silent or response file 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 Advanced database configuration option, then Oracle Universal Installer does not prompt you for database information. Instead, it installs the software and then runs Oracle Database Configuration Assistant in interactive mode. Using the screens in Oracle Database Configuration Assistant, you can either modify one of the preconfigured database types or customize the database.

Note: If you choose this method to create a database, then click **Help** on any of the Oracle Database Configuration Assistant 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, then you can use Oracle Database Configuration Assistant to create one after you have installed the software. For more information about using Oracle Database Configuration Assistant to create a database after installation, refer to the *Oracle Database 2 Day DBA* manual.

Database Storage Options

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

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

Note: Installing files on raw devices is no longer an option during installation. You must use a file system, or use Automatic Storage Management.

File System

If you choose the file system option, then Oracle Database Configuration Assistant creates the database files in a directory on a file system mounted on the 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, then Oracle recommends that you follow the Optimal Flexible Architecture (OFA) recommendations 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, then 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 also have the option to use the Direct NFS feature, which simplifies the administration of NFS configurations and also offers performance improvements.

See Also: ["Direct NFS Client"](#) for more information on the Direct NFS feature

If the NAS device is certified by Oracle, then you can store the database files on them.

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

See Also: *Oracle Database Administrator's Guide* for more information about Oracle-managed files

Automatic Storage Management

Automatic Storage Management is a high-performance storage management solution. For Oracle Database files, it simplifies the management of a dynamic database environment, such as creating and laying out databases and managing disk space.

Automatic Storage Management can be used with single database installations, multiple database installations, and in Oracle RAC environments. It can be used with databases created in Oracle Database 10g Release 1 (10.1.0.3 or later). However, Oracle Database 11g Release 2 (11.2) databases must use Automatic Storage Management from Oracle Database 11g Release 2 (11.2) or later. Automatic Storage Management is installed as part of the grid infrastructure installation. If you plan to use Automatic Storage Management, then you must install the grid infrastructure before installing your database. If you want to upgrade an existing Automatic Storage Management installation, then you must upgrade Automatic Storage Management by running an Oracle grid infrastructure upgrade.

See Also: [Chapter 3, "Oracle Grid Infrastructure"](#) for more information about installing the Oracle grid infrastructure software

Automatic Storage Management manages the storage of all database files, such as redo logs, control files, and data pump export files. Automatic Storage Management can manage the Oracle Database executable binary files as well as any other non-database file by creating a file system with Oracle Automatic Storage Management Cluster File System. Though Oracle Automatic Storage Management Cluster File System is cluster aware it works as a file system on a single instance database also.

See Also: *Oracle Database Storage Administrator's Guide* for more information about Oracle Automatic Storage Management Cluster File System

At a high level, implementing Automatic Storage Management involves allocating partitioned disks for Oracle Database with preferences for striping and mirroring.

Automatic Storage Management manages the disk space for you. This helps avoid the need for traditional disk management tools, such as Logical Volume Managers (LVM), file systems, and the numerous commands necessary to manage both. The synchronization between Automatic Storage Management and the database instance is handled by CSS.

The following are components of an Automatic Storage Management installation:

- [Automatic Storage Management Disk Groups](#)
- [Automatic Storage Management Instance](#)

Automatic Storage Management Disk Groups

A disk group is a set of disk devices that Automatic Storage Management manages as a single unit. Each disk device can be an individual physical disk, a multiple disk device, such as a RAID storage array or logical volume, or even a partition on a physical disk. However, in most cases, disk groups consist of one or more individual physical disks. To enable Automatic Storage Management to balance input-output operation and storage efficiently within the disk group, you must ensure that all devices in the disk group have similar, if not identical, storage capacity and performance.

You can set the redundancy and striping attributes of individual file types within a disk group by using Automatic Storage Management disk group templates. When you create a disk group, Automatic Storage Management creates a set of default templates for that disk group. Default template settings depend on the disk group type. For example, the default template for control files for a normal redundancy disk group sets three-way mirroring. All other file templates are two-way mirrored. For a high redundancy disk group, the default mirroring cannot be changed, which implies that all files are always three-way mirrored in a high redundancy disk group. You can modify the default templates to suit your site's needs. Refer to *Oracle Database Storage Administrator's Guide* for more information.

Automatic Storage Management spreads data evenly across all the devices in the disk group to optimize performance and utilization. You can add or remove disk devices from a disk group without shutting down the database. When you add or remove disks, Automatic Storage Management rebalances the files across the disk group. You can create multiple disk groups to handle specific tasks, such as backup and recovery operations, in addition to regular file storage activities.

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

Automatic Storage Management Instance

The Automatic Storage Management instance is a special Oracle instance that manages Automatic Storage Management disk groups. The Automatic Storage Management instance is created and started, if necessary, when you install Oracle grid infrastructure. Oracle recommends that you have the Automatic Storage Management instance in its own Oracle home. Oracle also recommends that you run this instance before you start a database instance that uses Automatic Storage Management.

For an Oracle Database installation, you only need one Automatic Storage Management instance, regardless of the number of database instances on the computer.

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:

- Deploy Oracle Enterprise Manager 10g centrally in the environment

To deploy Oracle Enterprise Manager centrally, you must install at least one Oracle Management Repository and one Oracle Management Service within the environment, then install an Oracle Enterprise Management Agent on every computer that you want to manage. You can then use a single HTML 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 on the Oracle Enterprise Manager Grid Control installation media, and also on the Oracle Technology Network Web site at:

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

- Deploy Oracle Enterprise Manager Database Control locally on the database system

Oracle Enterprise Manager Database Control software is installed by default 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, then you must either configure a separate Database Control for each database, or install Oracle Enterprise Manager 10g Grid Control.

Note: Refer to the *Oracle Enterprise Manager Concepts* manual and the *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* manual on the Oracle Enterprise Manager Grid Control installation media for more information about Oracle Enterprise Manager 10g.

This section contains the following topics:

- [Management Options for Preconfigured Databases](#)
- [Management Options for Custom Databases](#)
- [Features Provided by Oracle Enterprise Manager Database Control](#)

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 central database management

This option is available only if an Enterprise Manager Database Control Agent is installed on the system. When Oracle Universal Installer detects an Oracle Management Agent on the system, you can 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, then you must use Database Control to manage the database. However, if Oracle Management Agent is installed after Oracle Database, then you can use Grid Control to manage this database.

- Use Database Control for local 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, then Oracle Universal Installer runs Oracle Database Configuration Assistant in interactive mode. Using a screen in Oracle Database Configuration Assistant, 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.

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, then you can use Oracle Database Configuration Assistant 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 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 matrix

In addition, it provides you with automatic notification of security alerts, and 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, then you can optionally enable automated database backups that use the Oracle-suggested default backup strategy. 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.

This section contains the following topics:

- [Enabling Automated Backups](#)
- [Backup Job Default Settings](#)

See Also:

- *Oracle Database 2 Day DBA* for information about using Oracle Enterprise Manager Database Control to configure or customize automated backups or to recover a backed up database
- *Oracle Database Backup and Recovery User's Guide* for more detailed information about defining a backup strategy and backing up and recovering Oracle databases

Enabling Automated Backups

If you enable automated backups, then Oracle Enterprise Manager schedules a daily backup job that uses Oracle Recovery Manager (RMAN) to back up all of the database files to a disk storage area called the fast recovery area. The first time the backup job runs, it creates a full backup of the database. Subsequent backup jobs perform incremental backups, 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 fast recovery area

You can choose to use either a file system directory or an Automatic Storage Management disk group for the fast recovery area. To set the default values for fast recovery area and data file location, use Oracle base as the starting point.

- Default fast recovery area: `$ORACLE_BASE/recovery_area`
- Default data file location: `$ORACLE_BASE/oradata`

The default disk quota configured for the fast recovery area is 2 GB. For Automatic Storage Management 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 fast 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 `ORA_DBA` group). This user also must have Logon As A Batch Job privilege.

Backup Job Default Settings

If you enable automated backups after choosing one of the preconfigured databases during the installation, then 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 fast recovery area is 2 GB.

If you enable automated backups by using Oracle Database Configuration Assistant, either during or after the installation, then you can specify a different start time for the backup job and a different disk quota for the fast recovery area.

E-mail Notification Options

If you choose to use the Oracle Enterprise Manager Database Control during the installation, then you can configure Enterprise Manager to send an 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, then you must specify the following information:

- The host name of a Simple Mail Transfer 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.

Migration Consideration

Oracle Database 11g Release 2 (11.2) database for 32-bit Linux can be migrated to an Oracle Database 11g Release 2 (11.2) database for 64-bit Linux. Refer to "Database Migration from a 32-Bit Linux to 64-Bit Linux Computer" section in the *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for migration information.

Upgrade Considerations

For information about upgrading a earlier release of Oracle Database to Oracle Database 11g Release 2 (11.2), refer to *Oracle Database Upgrade Guide*. The following sections provide additional platform-specific upgrade information that you should review before upgrading an existing database:

- [Upgrading an Oracle Database Installed on Red Hat Enterprise Linux 2.1](#)
- [Automatic Storage Management Is Installed with Oracle Grid Infrastructure](#)
- [Daylight Savings Time Upgrade](#)

Upgrading an Oracle Database Installed on Red Hat Enterprise Linux 2.1

If you have the 8.1.7, 9.0.1, 9.2.0, or 10.1 release of Oracle Database installed on Red Hat Enterprise Linux 2.1, then you must first upgrade the operating system to Red Hat

Enterprise Linux 4 (update 3) before you upgrade the database. To do this, perform any one of the following procedures:

See Also: *Oracle Database Upgrade Guide*

- Upgrade the operating system. Then, upgrade the database either manually or by using Oracle Database Upgrade Assistant.
- Copy the database files. This procedure involves the following steps:
 1. Copy the database files from the computer running Red Hat Enterprise Linux 2.1 to the one running Red Hat Enterprise Linux 4.0.
 2. Re-create the control files on the computer running Red Hat Enterprise Linux 4.0.
 3. Manually upgrade the database.

Note: You cannot use Oracle Database Upgrade Assistant if you follow this method. However, this method lets you easily revert to the earlier database.

- Upgrade the database by using the Export/Import utilities.

Automatic Storage Management Is Installed with Oracle Grid Infrastructure

In previous releases, Oracle Automatic Storage Management (ASM) was installed as part of the Oracle Database installation. With Oracle Database 11g Release 2 (11.2), Automatic Storage Management is part of an Oracle grid infrastructure installation, either for a cluster, or for a standalone server.

If you want to upgrade an existing Automatic Storage Management installation, then you must upgrade Automatic Storage Management by running an Oracle grid infrastructure upgrade. If you do not have Automatic Storage Management installed and you want to use Automatic Storage Management as your storage option, then you must complete an Oracle grid infrastructure installation before you start your Oracle Database installation.

See Also: *Oracle Database Upgrade Guide*

Daylight Savings Time Upgrade

Refer to "[Daylight Savings Time Upgrade of Timestamp with Timezone Data Type](#)" for information about Daylight Savings Time Upgrade.

Oracle Database Preinstallation Requirements

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

Note:

- This guide contains information required to install Oracle Database 11g Release 2 (11.2) on various platforms of Linux. Ensure that you review information related to the platform on which you intend to install Oracle Database 11g.
 - If you want to use Automatic Storage Management or Oracle Restart, then you must first install Oracle grid infrastructure for a standalone server and then install Oracle Database.
-
-

- [Logging In to the System as root](#)
- [Checking the Hardware Requirements](#)
- [Checking the Software Requirements](#)
- [Preinstallation Requirement for Oracle Database Vault](#)
- [Installation Fixup Scripts](#)
- [Checking the Network Setup](#)
- [Creating Required Operating System Groups and Users](#)
- [Configure Oracle Installation Owner Shell Limits](#)
- [Configuring Kernel Parameters](#)
- [Identifying Required Software Directories](#)
- [Identifying or Creating an Oracle Base Directory](#)
- [Choosing a Storage Option for Oracle Database and Recovery Files](#)
- [Creating Directories for Oracle Database or Recovery Files](#)
- [Configuring Storage for Oracle Database Files Using Block Devices](#)
- [Configuring Disk Devices for Oracle Database](#)
- [Stopping Existing Oracle Processes](#)
- [Configuring the oracle User's Environment](#)

See Also: ■

- ["Requirements for Oracle Grid Infrastructure Installation"](#) for installation requirements for Oracle grid infrastructure
- "Pre-installation Requirements" section in *Oracle Configuration Manager Installation and Administration Guide* and *Oracle Configuration Manager Prerequisites* for preinstallation requirements for Oracle Configuration Manager

Logging In to the System as root

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

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

For more information about silent-mode installations, refer to [Appendix A](#).

- Following are the steps for 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, then enter the following command to enable the remote host to display X applications on the local X server:

```
$ xhost fully_qualified_remote_host_name
```

For example:

```
$ xhost somehost.us.example.com
```

3. If you are not installing the software on the local system, then use the `ssh`, `rlogin`, or `telnet` command to connect to the system where you want to install the software:

```
$ telnet fully_qualified_remote_host_name
```

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

```
$ sudo sh
password:
#
```

- Following are the steps for installing the software from a PC or other system with X server software:

Note: If necessary, refer to the 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, then enter the following command to switch user to `root`:

```
$ sudo sh
password:
#
```

Checking the Hardware Requirements

The system must meet the following minimum hardware requirements:

- [Memory Requirements](#)
- [System Architecture](#)
- [Disk Space Requirements](#)
- [Display Requirements](#)
- [Run Level Requirement](#)
- [Recommended Hardware Requirement for SQL Developer](#)

Memory Requirements

The following are the memory requirements for installing Oracle Database 11g Release 2:

- At least 1 GB of RAM

To determine the RAM size, enter the following command:

```
# grep MemTotal /proc/meminfo
```

If the size of the RAM is less than the required size, then you must install more memory before continuing.

- The following table describes the relationship between installed RAM and the configured swap space requirement:

Note: On Linux, the HugePages feature allocates non-swappable memory for large page tables using memory-mapped files. If you enable HugePages, then you should deduct the memory allocated to HugePages from the available RAM before calculating swap space.

RAM	Swap Space
Between 1 GB and 2 GB	1.5 times the size of RAM
Between 2 GB and 16 GB	Equal to the size of RAM
More than 16 GB	16 GB

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

```
# grep SwapTotal /proc/meminfo
```

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

To determine the available RAM and swap space, enter the following command:

```
# free
```

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

Automatic Memory Management

Starting with Oracle Database 11g, the Automatic Memory Management feature requires more shared memory (`/dev/shm`) and file descriptors. The size of the shared memory should be at least the greater of `MEMORY_MAX_TARGET` and `MEMORY_TARGET` for each Oracle instance on the computer. If `MEMORY_MAX_TARGET` or `MEMORY_TARGET` is set to a non zero value, and an incorrect size is assigned to the shared memory, it will result in an `ORA-00845` error at startup. The number of file descriptors for each Oracle instance should be at least `512 * PROCESSES`. Also, the limit of descriptors for each process should be at least 512. If file descriptors are not sized correctly, you will notice `ORA-27123` from various Oracle processes and potentially Linux Error `EMFILE` (Too many open files) errors in non-Oracle processes.

To determine the amount of shared memory available, enter the following command:

```
# df -h /dev/shm/
```

Note: `MEMORY_MAX_TARGET` and `MEMORY_TARGET` cannot be used when `LOCK_SGA` is enabled or with HugePages on Linux.

On the Initialization Parameters page, note that Memory Size (SGA and PGA), which sets the initialization parameter `MEMORY_TARGET` or `MEMORY_MAX_TARGET`. Note that the initialization parameters cannot be greater than the shared memory file system on the operating system. For example, if the shared memory file system allocation on your system is 1 GB, but you set Memory Size (`MEMORY_TARGET`) to 2 GB, then the following error messages are displayed during database startup:

```
ORA-00845: MEMORY_TARGET not supported on this system
ORA-01078: Failure in processing system parameters
```

In addition, if you click All Initialization Parameters and the global database name is longer than 8 characters, then the database name value (in the `DB_NAME` parameter) is truncated to the first eight characters, and the `DB_UNIQUE_NAME` parameter value is set to the global name.

System Architecture

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

```
# uname -m
```

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

Disk Space Requirements

The following are the disk space requirements for installing Oracle Database 11g Release 2 (11.2):

- At least 1 GB of disk space in the `/tmp` directory

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

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

If there is less than 1 GB of free disk space available in the `/tmp` directory, then complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to meet the disk space requirement.
- Set the `TMP` and `TMPDIR` environment variables when setting the `oracle` user's environment.

See Also: "[Configuring the oracle User's Environment](#)" for more information on setting `TMP` and `TMPDIR`

- Extend the file system that contains the `/tmp` directory. If necessary, contact the system administrator for information about extending file systems.
- To determine the amount of free disk space on the system, enter the following command:


```
# df -h
```
- The following tables describe the disk space requirements for software files, and data files for each installation type on Linux x86:

Installation Type	Requirement for Software Files (GB)
Enterprise Edition	3.95
Standard Edition	3.88

Installation Type	Disk Space for Data Files (GB)
Enterprise Edition	1.7
Standard Edition	1.5

- The following tables describe the disk space requirements for software files, and data files for each installation type on Linux x86-64:

Installation Type	Requirement for Software Files (GB)
Enterprise Edition	4.35
Standard Edition	3.73

Installation Type	Disk Space for Data Files (GB)
Enterprise Edition	1.68
Standard Edition	1.48

Additional disk space, either on a file system or on an Automatic Storage Management disk group is required for the fast recovery area if you choose to configure automated backups.

Display Requirements

The minimum display requirement for Oracle Database 11g Release 2 (11.2) is a resolution of 1024 x 768 or higher.

Run Level Requirement

Ensure that the system is started with run level 3 or run level 5.

Recommended Hardware Requirement for SQL Developer

The following table lists the recommended Memory and Display requirements for SQL Developer:

Resource	Recommended
CPU Type and Speed	Pentium IV 2 GHz or faster
Java SDK	JDK 5.0 Update 6 or later for Linux
Memory	1 GB RAM
Display	65536 colors, set to at least 1024 X 768 resolution

Checking the Software Requirements

Depending on the products that you intend to install, verify that the following softwares are installed on the system.

Note:

- This guide contains information required to install Oracle Database 11g Release 2 (11.2) on various platforms of Linux. Ensure that you review information related to the platform on which you intend to install Oracle Database 11g.
 - Oracle Universal Installer performs checks on the system to verify that it meets the listed requirements. To ensure that these checks pass, verify the requirements before you start Oracle Universal Installer.
-
-

- [Operating System Requirements](#)
- [Kernel Requirements](#)
- [Package Requirements](#)

- [Compiler Requirements](#)
- [Additional Software Requirements](#)

Operating System Requirements

The following are the operating system requirements for Oracle Database 11g Release 2 (11.2) for Linux x86, and Linux x86-64:

- Asianux 2.0
- Asianux 3.0
- Oracle Enterprise Linux 4.0 Update 7 or later
- Oracle Enterprise Linux 5.0
- Red Hat Enterprise Linux 4.0 Update 7 or later
- Red Hat Enterprise Linux 5.0
- SUSE Linux Enterprise Server 10.0
- SUSE Linux Enterprise Server 11.0

Starting with Oracle Database 11g Release 2 (11.2), the Security Enhanced Linux (SE Linux) feature is supported for Oracle Enterprise Linux 4.0, Red Hat Enterprise Linux 4.0, Oracle Enterprise Linux 5.0, and Red Hat Enterprise Linux 5.0.

Note: For Asianux Server, Oracle Enterprise Linux, and Red Hat Enterprise Linux, system requirements are identical by kernel version. Specifically:

Asianux 2, Oracle Enterprise Linux 4, and Red Hat Enterprise Linux 4 requirements are the same.

Asianux Server 3, Oracle Enterprise Linux 5, and Red Hat Enterprise Linux 5 requirements are the same.

See Also: ["Hardware and Software Certification"](#) for information about how to access the latest system requirements

To determine the distribution and version of Linux installed, enter the following command:

```
# cat /proc/version
```

Alternatively, you can also enter the following command on some distributions of Linux:

```
# lsb_release -id
```

Note: Only the distributions and versions listed in the earlier list are supported. Do not install the software on other versions of Linux.

Kernel Requirements

The following are the kernel requirements for Oracle Database 11g Release 2 (11.2):

- On Asianux 2.0, Oracle Enterprise Linux 4.0, and Red Hat Enterprise Linux 4.0:
2.6.9 or later
- On Asianux 3.0, Oracle Enterprise Linux 5.0, and Red Hat Enterprise Linux 5.0:
2.6.18 or later
- On SUSE Linux Enterprise Server 10:
2.6.16.21 or later
- On SUSE Linux Enterprise Server 11:
2.6.27.19 or later

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

```
# uname -r
```

The following is a sample output displayed by running this command on an Oracle Enterprise Linux 5.0 system:

```
2.6.18-128.el5PAE
```

In this example, the output shows the kernel version (2.6.18) and errata level (-128.el5PAE) on the system.

If the kernel version does not meet the requirement, then contact the operating system vendor for information about obtaining and installing kernel updates.

Package Requirements

The following are the list of packages required for Oracle Database 11g Release 2 (11.2):

Note:

- Oracle recommends that you install your Linux operating system with the default software packages (RPMs), unless you specifically intend to perform a minimal installation, and follow the directions for performing such an installation to ensure that you have all required packages for Oracle software.
 - Oracle recommends that you do not customize RPMs during a default operating system installation. A default installation includes most required packages, and will help you to limit manual checks of package dependencies.
 - If you did not perform a default Linux installation, you intend to use LDAP, and you want to use the scripts `odisrvreg`, `oidca`, or `schemasync`, then install the Korn shell RPM for the Linux distribution.
 - You must install the packages (or later versions) listed in the following table. Also, ensure that the list of RPMs and all of the prerequisites for these RPMs are installed.
-
-

On Linux x86:

Operating System	Requirement
Asianux 2.0, Oracle Enterprise Linux 4.0, and Red Hat Enterprise Linux 4.0	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.15.92.0.2 compat-libstdc++-33.2.3 elfutils-libelf-0.97 elfutils-libelf-devel-0.97 gcc-3.4.6 gcc-c++-3.4.6 glibc-2.3.4-2.41 glibc-common-2.3.4 glibc-devel-2.3.4 glibc-headers-2.3.4 libaio-devel-0.3.105 libaio-0.3.105 libgcc-3.4.6 libstdc++-3.4.6 libstdc++-devel-3.4.6 make-3.80 pdksh-5.2.14 sysstat-5.0.5 unixODBC-2.2.11 unixODBC-devel-2.2.11 </pre>
Asianux 3.0, Oracle Enterprise Linux 5.0, and Red Hat Enterprise Linux 5.0	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.17.50.0.6 compat-libstdc++-33-3.2.3 elfutils-libelf-0.125 elfutils-libelf-devel-0.125 elfutils-libelf-devel-static-0.125 gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-24 glibc-common-2.5 glibc-devel-2.5 glibc-headers-2.5 kernel-headers-2.6.18 ksh-20060214 libaio-0.3.106 libaio-devel-0.3.106 libgcc-4.1.2 libgomp-4.1.2 libstdc++-4.1.2 libstdc++-devel-4.1.2 make-3.81 sysstat-7.0.2 unixODBC-2.2.11 unixODBC-devel-2.2.11 </pre>

Operating System	Requirement
SUSE Linux Enterprise Server 10	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.16.91.0.5 compat-libstdc++-5.0.7 gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-24 glibc-devel-2.4 ksh-93r-12.9 libaio-0.3.104 libaio-devel-0.3.104 libelf-0.8.5 libgcc-4.1.2 libstdc++-4.1.2 libstdc++-devel-4.1.2 make-3.80 sysstat-8.0.4 </pre>
SUSE Linux Enterprise Server 11	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.19 gcc-4.3 gcc-c++-4.3 glibc-2.9 glibc-devel-2.9 ksh-93t libstdc++33-3.3.3 libstdc++43-4.3.3_20081022 libstdc++43-devel-4.3.3_20081022 libaio-0.3.104 libaio-devel-0.3.104 libgcc43-4.3.3_20081022 libstdc++-devel-4.3 make-3.81 sysstat-8.1.5 </pre>

On Linux x86-64:

Operating System	Requirement
Asianux 2.0, Oracle Enterprise Linux 4.0, and Red Hat Enterprise Linux 4.0	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.15.92.0.2 compat-libstdc++-33-3.2.3 compat-libstdc++-33-3.2.3 (32 bit) elfutils-libelf-0.97 elfutils-libelf-devel-0.97 expat-1.95.7 gcc-3.4.6 gcc-c++-3.4.6 glibc-2.3.4-2.41 glibc-2.3.4-2.41 (32 bit) glibc-common-2.3.4 glibc-devel-2.3.4 glibc-headers-2.3.4 libaio-0.3.105 libaio-0.3.105 (32 bit) libaio-devel-0.3.105 libaio-devel-0.3.105 (32 bit) libgcc-3.4.6 libgcc-3.4.6 (32-bit) libstdc++-3.4.6 libstdc++-3.4.6 (32 bit) libstdc++-devel 3.4.6 make-3.80 pdksh-5.2.14 sysstat-5.0.5 unixODBC-2.2.11 unixODBC-2.2.11 (32 bit) unixODBC-devel-2.2.11 unixODBC-devel-2.2.11 (32 bit) </pre>

Operating System	Requirement
Asianux 3.0, Oracle Enterprise Linux 5.0, and Red Hat Enterprise Linux 5.0	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.17.50.0.6 compat-libstdc++-33-3.2.3 compat-libstdc++-33-3.2.3 (32 bit) elfutils-libelf-0.125 elfutils-libelf-devel-0.125 gcc-4.1.2 gcc-c++-4.1.2 glibc-2.5-24 glibc-2.5-24 (32 bit) glibc-common-2.5 glibc-devel-2.5 glibc-devel-2.5 (32 bit) glibc-headers-2.5 ksh-20060214 libaio-0.3.106 libaio-0.3.106 (32 bit) libaio-devel-0.3.106 libaio-devel-0.3.106 (32 bit) libgcc-4.1.2 libgcc-4.1.2 (32 bit) libstdc++-4.1.2 libstdc++-4.1.2 (32 bit) libstdc++-devel 4.1.2 make-3.81 sysstat-7.0.2 unixODBC-2.2.11 unixODBC-2.2.11 (32 bit) unixODBC-devel-2.2.11 unixODBC-devel-2.2.11 (32 bit) </pre>
SUSE Linux Enterprise Server 10	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.16.91.0.5 compat-libstdc++-5.0.7 gcc-4.1.0 gcc-c++-4.1.2 glibc-2.5-24 glibc-devel-2.4 glibc-devel-32bit-2.4 ksh-93r-12.9 libaio-0.3.104 libaio-32bit-0.3.104 libaio-devel-0.3.104 libaio-devel-32bit-0.3.104 libelf-0.8.5 libgcc-4.1.2 libstdc++-4.1.2 libstdc++-devel-4.1.2 make-3.80 sysstat-8.0.4 </pre>

Operating System	Requirement
SUSE Linux Enterprise Server 11	<p>The following packages (or later versions) must be installed:</p> <pre> binutils-2.19 gcc-4.3 gcc-32bit-4.3 gcc-c++-4.3 glibc-2.9 glibc-32bit-2.9 glibc-devel-2.9 glibc-devel-32bit-2.9 ksh-93t libaio-0.3.104 libaio-32bit-0.3.104 libaio-devel-0.3.104 libaio-devel-32bit-0.3.104 libstdc++33-3.3.3 libstdc++33-32bit-3.3.3 libstdc++43-4.3.3_20081022 libstdc++43-32bit-4.3.3_20081022 libstdc++43-devel-4.3.3_20081022 libstdc++43-devel-32bit-4.3.3_20081022 libgcc43-4.3.3_20081022 libstdc++-devel-4.3 make-3.81 sysstat-8.1.5 </pre>

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

```
# rpm -q package_name
```

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

Compiler Requirements

Intel C++ Compiler 10.1 or later and the version of GNU C and C++ compilers listed under "[Package Requirements](#)" on page 8 are supported with Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, and Oracle XML Developer's Kit (XDK) for Oracle Database 11g Release 2.

Note: Intel Compiler v10.1 can be used only with gcc 3.4.5 or gcc 4.0 or gcc 4.1 standard template libraries to build Oracle C++ Call Interface (OCCI) applications.

Oracle XML Developer's Kit is supported with the same compilers as OCCI.

Additional Software Requirements

Depending on the components you want to use, you must ensure that the following software are installed:

- [Oracle ODBC Drivers](#)
- [Oracle JDBC/OCI Drivers](#)

- [Oracle Messaging Gateway](#)
- [Browser Requirements](#)

Oracle ODBC Drivers

If you intend to use ODBC, then you should install the most recent ODBC Driver Manager for Linux. You can download and install the Driver Manager from the following URL:

<http://www.unixodbc.org>

Linux RPMs are also available on this Web site.

- On Linux x86
 - To use ODBC, you must also install the following additional 32-bit ODBC RPMs, depending on your operating system:
 - On Asianux 2, Oracle Enterprise Linux 4, and Red Hat Enterprise Linux 4:
unixODBC-2.2.11 (32 bit) or later
unixODBC-devel-2.2.11 (32 bit) or later
 - On Asianux 3, Oracle Enterprise Linux 5, and Red Hat Enterprise Linux 5:
unixODBC-2.2.11 (32 bit) or later
unixODBC-devel-2.2.11 (32 bit) or later
 - On SUSE 10:
unixODBC-32bit-2.2.11 (32 bit) or later
unixODBC-devel-32bit-2.2.11 (32 bit) or later
- On Linux x86-64
 - To use ODBC, you must also install the following additional ODBC RPMs, depending on your operating system:
 - On Asianux 2, Oracle Enterprise Linux 4, and Red Hat Enterprise Linux 4:
unixODBC-2.2.11 (32 bit) or later
unixODBC-devel-2.2.11 (64 bit) or later
unixODBC-2.2.11 (64 bit) or later
 - On Asianux 3, Oracle Enterprise Linux 5, and Red Hat Enterprise Linux 5:
unixODBC-2.2.11 (32 bit) or later
unixODBC-devel-2.2.11 (64 bit) or later
unixODBC-2.2.11 (64 bit) or later
 - On SUSE 10:
unixODBC-32bit-2.2.11 (32 bit) or later
unixODBC-2.2.11 (64 bit) or later
unixODBC-devel-2.2.11 (64 bit) or later

Oracle JDBC/OCI Drivers

You can use Sun JDK 1.5.0-06 with the JNDI extension with the Oracle Java Database Connectivity and Oracle Call Interface drivers. However, these are not mandatory for the database installation.

Oracle Messaging Gateway

Oracle Messaging Gateway supports the integration of Oracle Streams Advanced Queuing (AQ) with the following software:

- IBM WebSphere MQ V6.0, client and server, with corrective service diskette 5 (CSD05) or later:

```
MQSeriesClient
MQSeriesServer
MQSeriesRuntime
```

- TIBCO Rendezvous 7.3

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

<http://www-306.ibm.com/software/integration/wmq/support>

Browser Requirements

Web browsers must support Java Script, and the HTML 4.0 and CSS 1.0 standards. The following browsers meet these requirements for Oracle Enterprise Manager Database Control:

- Netscape Navigator 8.1
- Netscape Navigator 9.0
- Microsoft Internet Explorer 6.0 SP2
- Microsoft Internet Explorer 7.0
- Microsoft Internet Explorer 7.0 SP1
- Microsoft Internet Explorer 8.0
- Firefox 2.0
- Firefox 3.0.7
- Safari 3.0.4
- Safari 3.1
- Safari 3.2

See Also: Chapter 2, "Oracle Application Express Installation Requirements" and "Recommended Pre-installation Tasks" in *Oracle Application Express Installation Guide*

Preinstallation Requirement for Oracle Database Vault

If you want to install Oracle Database Vault, then set the `DB_BLOCK_SIZE` initialization parameter to 4096 or larger. If the value is less than 4096, then you cannot change it. The only way to change the `DB_BLOCK_SIZE` value is by recreating the database.

See Also: "Specifying Database Block Sizes" in *Oracle Database Administrator's Guide*

Installation Fixup Scripts

During installation, for certain prerequisite check failures, you can click **Fix & Check Again** to generate a fixup script (`runfixup.sh`). You can run this script as a `root` user to complete the required preinstallation steps.

The fixup script does the following:

- Checks and sets kernel parameters to values required for successful installation, including:
 - Shared memory parameters
 - Semaphore parameters
 - Open file descriptor and UDP send/receive parameters
- Sets permissions on the Oracle Inventory directory.
- Reconfigures primary and secondary group memberships for the installation owner, if necessary, for the Oracle Inventory directory, and for the operating system privileges groups.
- Sets up virtual IP and private IP addresses in `/etc/hosts`.
- Sets shell limits to required values, if necessary.
- Installs the Cluster Verification Utility packages (`cvuqdisk rpm`).

Oracle recommends that you do not modify the contents of the generated fixup script.

Note: Using fixup scripts will not ensure that all the prerequisites for installing Oracle Database are satisfied. You must still verify that all the preinstallation requirements are met to ensure a successful installation.

Checking the Network Setup

Typically, the computer on which you want to install Oracle Database is connected to the network. The computer has local storage, to store the Oracle Database installation. It also contains a display monitor, and DVD drive. This section describes how to install Oracle Database on computers that do not meet the typical scenario. It covers the following cases:

- [Installing on DHCP Computers](#)
- [Installing on Multihomed Computers](#)
- [Installing on Computers with Multiple Aliases](#)
- [Installing on Non-Networked Computers](#)

Installing on DHCP Computers

Dynamic Host Configuration Protocol (DHCP) assigns dynamic IP addresses on a network. Dynamic addressing enables a computer to have a different IP address each time it connects to the network. In some cases, the IP address can change while the computer is still connected. You can have a mixture of static and dynamic IP addressing in a DHCP system.

In a DHCP setup, the software tracks IP addresses, which simplifies network administration. This lets you add a new computer to the network without having to manually assign a unique IP address to the newly added computer.

Installing on Multihomed Computers

You can install Oracle Database on a multihomed computer. A multihomed computer is associated with multiple IP addresses. This is typically achieved by having multiple network cards on the computer. Each IP address is associated with a host name. In addition, you can set up aliases for the host name. By default, Oracle Universal Installer uses the `ORACLE_HOSTNAME` environment variable setting to find the host name. If `ORACLE_HOSTNAME` is not set and you are installing on a computer that has multiple network cards, then Oracle Universal Installer determines the host name by using the first entry in the `/etc/hosts` file.

Clients must be able to access the computer either by using this host name or by using aliases for this host name. To verify this, ping the host name from the client computers using the short name (host name only) and the full name (host name and domain name). Both tests must be successful.

Setting the `ORACLE_HOSTNAME` Environment Variable

Use the following procedure to set the `ORACLE_HOSTNAME` environment variable. For example, if the fully qualified host name is `somehost.us.example.com`, then enter one of the following commands:

In Bourne, Bash, or Korn shell:

```
$ ORACLE_HOSTNAME=somehost.us.example.com
$ export ORACLE_HOSTNAME
```

In C shell:

```
% setenv ORACLE_HOSTNAME somehost.us.example.com
```

Installing on Computers with Multiple Aliases

A computer with multiple aliases is registered with the naming service under a single IP but with multiple aliases. The naming service resolves any of those aliases to the same computer. Before installing Oracle Database on such a computer, set the `ORACLE_HOSTNAME` environment variable to the computer whose host name you want to use.

Installing on Non-Networked Computers

You can install Oracle Database on a non-networked computer. If the computer, such as a laptop, is configured for DHCP and you plan to connect the computer to the network after the Oracle Database installation, then use the `ping` command on the computer on which you want to install the database to check if the computer can connect to itself. Perform this step by first using only the host name and then using the fully qualified name, which should be in the `/etc/hosts` file.

Note: When you run the `ping` command on the computer itself, the `ping` command should return the IP address of the computer.

If the `ping` command fails, then contact the system administrator.

Connecting the Computer to the Network after Installation

If you connect the computer to a network after installation, then the Oracle Database instance on the computer can work with other instances on the network. The computer can use a static IP or DHCP, depending on the network to which you are connected.

Creating Required Operating System Groups and Users

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 operating system groups and users.

If you prefer to allocate operating system user privileges so that you can use one administrative user and one group for operating system authentication for all administrative privileges, then you can use the `oracle` user as the installation owner, and use one group as the primary group for any user requiring administrative privileges for Automatic Storage Management, and Oracle Database administration. This group must also be the Oracle Inventory group. To simplify using the defaults for Oracle tools the group name should be `oinstall`.

You can also create custom configuration groups and users based on job role separation. A custom configuration is a configuration with groups and users that divide access privileges granted by membership in separate operating system groups and users. You can create a single user (for example, `oracle`) to own both Oracle database, and Oracle grid infrastructure installations. Alternatively, you can create a separate user (for example, `grid`) to own the Oracle grid infrastructure installation.

- [Creating Custom Configuration Groups and Users for Job Roles](#)
- [Creating Database Operating System Groups and Users with Job Role Separation](#)

Note: In Oracle documentation, a user created to own only Oracle grid infrastructure software installations is called the `grid` user. A user created to own either all Oracle installations, or only Oracle database installations, is called the `oracle` user.

Creating Custom Configuration Groups and Users for Job Roles

This section provides an overview of how to create users and groups to divide access privileges by job roles. Log in as `root` to create these groups and users.

- [Users for Oracle Installations with Job Role Separation](#)
- [Database Groups for Job Role Installations](#)
- [Grid Infrastructure Groups for Job Role Installations](#)

Users for Oracle Installations with Job Role Separation

Oracle recommends that you create one software owner to own each Oracle software installation (typically, `oracle`, for the database software and `grid` for the Oracle Restart owner user). You must create at least one software owner the first time you install Oracle software on the system.

Oracle software owners must have the Oracle Inventory group as their primary group, so that each Oracle software installation owner can write to the Central Inventory. The database software owner (typically, `oracle`) must also have the OSDBA group of the grid infrastructure home so that database instances can log on to Automatic Storage

Management, and (if you create it) the OSOPER group as secondary groups. In Oracle documentation, the Oracle software owner users are referred to as `oracle` users.

See Also: *Oracle Database Administrator's Guide* for more information about the OSDBA, OSASM and OSOPER groups, and the SYSDBA, SYSASM and SYSOPER privileges

Database Groups for Job Role Installations

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

- The OSDBA group (typically, `dba`)

You must create this group the first time you install Oracle Database software on the system. This group identifies operating system user accounts that have database administrative privileges (the SYSDBA privilege). The name used for this group in Oracle code examples is `dba`.

- The OSOPER group for Oracle Database (typically, `oper`)

This is an optional group. Create this group if you want a separate group of operating system users to have a limited set of database administrative privileges (the SYSOPER privilege). This group cannot directly connect as SYSOPER, unless explicitly granted. However, they will have the privileges granted by the SYSOPER privilege. By default, even members of the OSDBA group have all privileges granted by the SYSOPER privilege.

Oracle Universal Installer prompts you to specify the name of this group. The usual name chosen for this group is `oper`.

Grid Infrastructure Groups for Job Role Installations

The following operating system groups and user are required if you are installing Oracle grid infrastructure:

Note: You can designate a unique group, separate from database administrator groups, or you can use the same group as the OSASM and OSDBA group, to grant system privileges to administer both the ASM instances and Oracle Database instances.

- The OSDBA group for Automatic Storage Management (typically, `asmdba`)

The OSDBA group for Automatic Storage Management can be the same group used as the OSDBA group for the database, or you can create a separate OSDBA group for Automatic Storage Management (typically, `asmdba`) to provide administrative access to Automatic Storage Management instances.

The Oracle grid infrastructure software owner (typically, `grid`) must be a member of the OSDBA group. Membership in the OSDBA group allows access to the files managed by Automatic Storage Management. If you have a separate OSDBA group for Automatic Storage Management, then the Oracle Restart software owner must be a member of the OSDBA group for each database and the OSDBA group for Automatic Storage Management.

- The OSASM group for Automatic Storage Management (typically, `asmadmin`)

SYSASM privileges for Automatic Storage Management files provide administrator privileges for storage file. In Oracle documentation, the operating system group

whose members are granted `SYSASM` privileges is called the OSASM group, and in command lines, is referred to as `asmadmin`. Automatic Storage Management can support multiple databases.

Members of the OSASM group can use SQL to connect to an Automatic Storage Management instance as `SYSASM` using operating system authentication. The `SYSASM` privileges permit mounting and dismounting disk groups, and other storage administration tasks. `SYSASM` privileges provide no access privileges on an RDBMS instance.

If you do not designate a separate group as the OSASM group, then the OSDBA group you define is also by default the OSASM group.

- The OSOPER group for Automatic Storage Management (typically, `asmoper`)

This is an optional group. Create this group if you want a separate group of operating system users to have a limited set of Oracle ASM instance administrative privileges (the `SYSOPER` for ASM privilege), including starting up and stopping the Oracle ASM instance. By default, members of the OSASM group also have all privileges granted by the `SYSOPER` for ASM privilege.

If you want to have an OSOPER for ASM group, then the grid infrastructure owner must be a member of this group.

Creating Database Operating System Groups and Users with Job Role Separation

The following sections describe how to create the required operating system user and groups:

- [Creating the Oracle Inventory Group](#)
- [Creating the OSDBA Group for Database Installations](#)
- [Creating an OSOPER Group for Database Installation](#)
- [Creating the OSASM Group for Automatic Storage Management](#)
- [Creating the OSDBA Group for Oracle Grid Infrastructure](#)
- [Creating the OSOPER Group for Automatic Storage Management](#)
- [Creating the Oracle Software Owner User](#)

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

Oracle recommends that you do not use the UID and GID defaults on each node, as group and user IDs likely will be different on each node. Instead, provide common assigned group and user IDs, and confirm that they are unused on any node before you create or modify groups and users.

Creating the Oracle Inventory Group

When you install Oracle software on the system for the first time, Oracle Universal Installer creates the `oraInst.loc` file. This file identifies the name of the Oracle Inventory group (typically, `oinstall`), and the path of the Oracle Inventory directory.

You can configure one group to be the access control group for the Oracle Inventory, for database administrators (OSDBA), and for all other access control groups used by Oracle software for operating system authentication. However, this group then must be the primary group for all users granted administrative privileges.

Log in as `root`, and use the following instructions to locate or create the Oracle Inventory group and a software owner:

- [Determining Whether the Oracle Inventory Group Exists](#)
- [Creating the Oracle Inventory Group](#)

Determining Whether the Oracle Inventory Group Exists

An `oraInst.loc` file has contents similar to the following:

```
inventory_loc=central_inventory_location
inst_group=group
```

In the preceding example, `central_inventory_location` is the location of the Oracle Central Inventory, and `group` is the name of the group that has permissions to write to the central inventory.

If you have an existing Oracle Inventory, then ensure that you use the same Oracle Inventory for all Oracle software installations, and ensure that all Oracle software users you intend to use for installation have permissions to write to this directory.

To determine if the Oracle Inventory group exist, enter the following command:

```
# grep oinstall /etc/group
```

To determine whether the `oraInst.loc` file exists, enter the following command:

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

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

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

In the previous output example:

- The `inventory_loc` group shows the location of the Oracle Inventory
- The `inst_group` parameter shows the name of the Oracle Inventory group (in this example, `oinstall`).

Creating the Oracle Inventory Group

If the `oraInst.loc` file does not exist, then create the Oracle Inventory group by entering the following command:

```
# /usr/sbin/groupadd oinstall
```

Creating the OSDBA Group for Database Installations

You must create an OSDBA group in the following circumstances:

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

If the OSDBA group does not exist or if you require a new OSDBA group, then create it as follows. In the following procedure, use the group name `dba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 502 dba
```

Creating an OSOPER Group for Database Installation

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

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

If you require a new OSOPER group (typically, `oper`), then create it as follows. In the following, use the group name `oper` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 505 oper
```

Creating the OSASM Group for Automatic Storage Management

If the OSASM group does not exist or if you require a new OSASM group, then create it as follows. In the following procedure, use the group name `asmadmin` unless a group with that name already exists:

```
# /usr/sbin/groupadd -504 asmadmin
```

Creating the OSDBA Group for Oracle Grid Infrastructure

If you require a new OSDBA group for Oracle grid infrastructure, then create it as follows. In the following procedure, use the group name `asmdba` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 506 asmdba
```

Creating the OSOPER Group for Automatic Storage Management

If you require an OSOPER group, then create it as follows. In the following procedure, use the group name `asmoper` unless a group with that name already exists:

```
# /usr/sbin/groupadd -g 505 asmoper
```

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 operating system user, with different group membership, to give database administrative privileges to those groups in a new Oracle Database installation.
- If you have created an Oracle software owner for Oracle grid infrastructure, such as `grid`, and you want to create a separate Oracle software owner for Oracle Database software, such as `oracle`.

Determining if an Oracle Software Owner User Exists To determine whether an Oracle software owner user named `oracle`, or `grid` exists, enter a command similar to the following (in this case, to determine if `oracle` exists):

```
# id oracle
```

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

```
uid=501(oracle) gid=501(oinstall) groups=502(dba),503(oper)
```

Determine whether you want to use the existing user, or create another user. If you want to use the existing user, then ensure that the user's primary group is the Oracle Inventory group (`oinstall`) and that it is a member of the appropriate OSDBA and OSOPER groups. Refer to one of the following sections for more information:

- [Creating an Oracle Software Owner User](#)
- [Modifying an Existing Oracle Software Owner User](#)

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

Creating an Oracle Software Owner User If the Oracle software owner user does not exist, or if you require a new Oracle software owner user, such as `oracle` or `grid`, then create it as described in this section (in this case to create the `oracle` user).

In the following procedure, use the user name `oracle` unless a user with that name already exists:

1. To create an `oracle` user, enter a command similar to the following:

```
# /usr/sbin/useradd -u 502 -g oinstall -G dba,asmdba,[oper] oracle
```

In the preceding command:

- The `-u` option specifies the user ID. Using this command flag is optional, as you can allow the system to provide you with an automatically generated user ID number. However, you must make note of the `oracle` user ID number, as you require it later during preinstallation.
 - 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 and ASMDBA groups. For example: `dba`, or `asmdba`, `oper`
2. Set the password of the `oracle` user:

```
# passwd oracle
```

Modifying an Existing Oracle Software Owner User If the `oracle` user exists, but its primary group is not `oinstall`, or it is not a member of the appropriate OSDBA or OSOPER groups, then enter a command similar to the following to modify it. Specify the primary group using the `-g` option and any required secondary group using the `-G` option:

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

Configure Oracle Installation Owner Shell Limits

For information, review "[Configuring the oracle User's Environment](#)" on page 2-38

To improve the performance of the software, you must increase the following shell limits for the `oracle` and `grid` user:

Shell Limit	Item in limits.conf	Hard Limit
Maximum number of open file descriptors	<code>nofile</code>	65536
Maximum number of processes available to a single user	<code>nproc</code>	16384
Maximum size of the stack segment of the process	<code>stack</code>	10240

To increase the shell limits:

Note: Depending on your environment, make the following changes for all Oracle installation owners. The steps below only show the user `oracle`.

1. Add the following lines to the `/etc/security/limits.conf` file: (the following example shows the software account owner `oracle`):

```
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    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 file on SUSE Linux Enterprise Server systems `/etc/profile.local`):

```
if [ $USER = "oracle" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
        ulimit -p 16384
        ulimit -n 65536
    else
        ulimit -u 16384 -n 65536
    fi
umask 022
fi
```

- For the C shell (`csh` or `tcsh`), add the following lines to the `/etc/csh.login` file (or the file on SUSE Linux Enterprise Server systems `/etc/csh.login.local`):

```
if ( $USER == "oracle" ) then
    limit maxproc 16384
    limit descriptors 65536
endif
```

Configuring Kernel Parameters

During installation, you can generate and run the Fixup script to check and set the kernel parameter values required for successful installation of the database. This script updates required kernel packages if necessary to minimum values.

If you cannot use the Fixup scripts, then verify that the kernel parameters shown in the following table are set to values greater than or equal to the minimum value shown. The procedure following the table describes how to verify and set the values manually.

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

Parameter	Minimum Value	File
semmsl	250	/proc/sys/kernel/sem
semmns	32000	
semopm	100	
semmni	128	
shmall	2097152	/proc/sys/kernel/shmall
shmmax	Either 4 GB - 1 byte, or half the size of physical memory (in bytes), whichever is lower. Default: 536870912	/proc/sys/kernel/shmmax
shmmni	4096	/proc/sys/kernel/shmmni
file-max	6815744	/proc/sys/fs/file-max
aio-max-nr	Maximum: 1048576 Note: This value limits concurrent outstanding requests and should be set to avoid I/O subsystem failures.	/proc/sys/fs/aio-max-nr
ip_local_port_range	Minimum: 9000 Maximum: 65500	/proc/sys/net/ipv4/ip_local_port_range
rmem_default	262144	/proc/sys/net/core/rmem_default
rmem_max	4194304	/proc/sys/net/core/rmem_max
wmem_default	262144	/proc/sys/net/core/wmem_default
wmem_max	1048576	/proc/sys/net/core/wmem_max

Note: If the current value for any parameter is greater than the value listed in this table, then the Fixup scripts do not change the value of that parameter.

See Also: ■

- ["Guidelines for Setting Semaphore Parameters"](#)
- ["Installation Fixup Scripts"](#)

Displaying and Changing Kernel Parameter Values

Enter the commands shown in the following table to display the current values of the kernel parameters, make a note of these values and identify any values that you must change:

Parameter	Command
semmsl, semmns, semopm, and semmni	# /sbin/sysctl -a grep sem This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and shmmni	# /sbin/sysctl -a grep shm This command displays the details of the shared memory segment sizes.
file-max	# /sbin/sysctl -a grep file-max This command displays the maximum number of file handles.
ip_local_port_ range	# /sbin/sysctl -a grep ip_local_port_range This command displays a range of port numbers.
rmem_default	# /sbin/sysctl -a grep rmem_default
rmem_max	# /sbin/sysctl -a grep rmem_max
wmem_default	# /sbin/sysctl -a grep wmem_default
wmem_max	# /sbin/sysctl -a grep wmem_max
aio-max-nr	# /sbin/sysctl -a grep aio-max-nr

If the value of any kernel parameter is different from the minimum value, then perform the following:

1. 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 minimum value, then specify the larger value.

```
fs.aio-max-nr = 1048576
fs.file-max = 6815744
kernel.shmall = 2097152
kernel.shmmax = 536870912
kernel.shmmni = 4096
```



```
kernel.sem = 250 32000 100 128
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048586
```

Note: The minimum value required for `shmmax` is 0.5 GB. However, Oracle recommends that you set the value of `shmmax` to 2.0 GB for optimum performance of the system.

By specifying the values in the `/etc/sysctl.conf` file, they persist when you restart the system. However, on SUSE Linux Enterprise Server systems, enter the following command to ensure that the system reads the `/etc/sysctl.conf` file when it restarts:

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

2. 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.

3. Enter the command `/sbin/sysctl -a` to confirm that the values are set correctly.
4. On SUSE systems only, enter the following command to cause the system to read the `/etc/sysctl.conf` file when it restarts:

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

5. On SUSE systems only, you must enter the GID of the `oinstall` group as the value for the parameter `/proc/sys/vm/hugetlb_shm_group`. Doing this grants members of `oinstall` a group permission to create shared memory segments.

For example, where the `oinstall` group GID is 501:

```
# echo 501 > /proc/sys/vm/hugetlb_shm_group
```

After running this command, use `vi` to add the following text to `/etc/sysctl.conf`, and enable the `boot.sysctl` script to run on system restart:

```
vm.hugetlb_shm_group=501
```

Note: Only one group can be defined as the `vm.hugetlb_shm_group`.

6. After updating the values of kernel parameters in the `/etc/sysctl.conf` file, either restart the computer, or run the command `sysctl -p` to make the changes in the `/etc/sysctl.conf` file available in the active kernel memory.

Identifying Required Software Directories

You must identify or create the following directories for the Oracle software:

- [Oracle Base Directory](#)
- [Oracle Inventory Directory](#)
- [Oracle Home Directory](#)

Note: Ensure that the paths you select for Oracle software, such as Oracle home paths and the Oracle base path, use only ASCII characters. Because installation owner names are used by default for some path, this ASCII character restriction applies to user names, file names, and directory names.

Oracle Base Directory

The Oracle base directory is a top-level directory for Oracle software installations. On Linux 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/software_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 can choose another mount point directory, such as `/oracle` or `/opt/oracle`.

- *software_owner* is the operating system user name of the software owner installing the Oracle software, for example `oracle`, or `grid`.

Note: If you start a database instance using `spfile` with `ORACLE_BASE` environment variable set, then its value is automatically stored in `spfile`. If you unset `ORACLE_BASE` environment variable subsequently and start the instance afresh, then database uses the value of Oracle base stored in `spfile`.

You must specify the Oracle base folder that contains all Oracle products.

Note: If you have an existing Oracle base, then you can select it from the Use existing list. By default, the list contains the existing value for Oracle base preselected. Refer to "[Installing the Oracle Database Software](#)" on page 2-6 for further information.

If you do not have an Oracle base, then you can create one by editing the text in the list box.

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 operating system users install Oracle software on the same system, then each user must create a

separate Oracle base directory. The following are examples of Oracle base directories that can exist on the same system:

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

Refer to ["Creating an Oracle Base Directory"](#) for information on creating an Oracle base directory.

Oracle Inventory Directory

The Oracle Inventory directory (`oraInventory`) stores an inventory of all software installed on the system. It is required and shared by all Oracle software installations on a single system. If you have an existing Oracle Inventory path, then Oracle Universal Installer continues to use that Oracle Inventory.

The first time you install Oracle software on a system, Oracle Universal Installer checks if you have created an OFA-compliant directory structure with the format `u[01-09]/app`, such as `/u01/app`, and that the user running the installation has permissions to write to that path. If this is true, then Oracle Universal Installer creates the Oracle Inventory directory similar to `/u[01-09]/app/oraInventory`. For example:

```
/u01/app/oraInventory
```

If you have set the environment variable `ORACLE_BASE` for the `oracle` user, then Oracle Universal Installer creates the Oracle Inventory directory similar to `$ORACLE_BASE/./oraInventory`. For example, if `ORACLE_BASE` is set to `/opt/oracle/11`, then the Oracle Inventory directory is created similar to `/opt/oracle/oraInventory`.

If you have neither created an OFA-compliant path nor set `ORACLE_BASE`, then the Oracle Inventory directory is placed in the home directory of the user that is performing the installation. For example:

```
/home/oracle/oraInventory
```

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

Note:

- All Oracle software installations rely on this directory. Ensure that you back it up regularly.
 - Do not delete this directory unless you have completely removed all Oracle software from the system.
 - By default, the Oracle Inventory directory is not installed under the Oracle Base directory. This is because all Oracle software installations share a common Oracle Inventory, so there is only one Oracle Inventory for all users. Whereas, there is a separate Oracle Base for each user.
-
-

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 Oracle Universal 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/11.2.0/dbhome_1
```

Oracle Universal 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.

Note: During installation, you must not specify an existing directory that has predefined permissions applied to it as the Oracle home directory. If you do, then you may experience installation failure due to file and group ownership permission errors.

Identifying or Creating an Oracle Base Directory

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

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

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

Identifying an Existing Oracle Base Directory

Existing Oracle base directories may not have paths that comply with OFA (Optimal Flexible Architecture) guidelines. However, if you identify an existing Oracle Inventory directory or existing Oracle home directories, then you can usually identify the Oracle base directories, as follows:

- Identifying an existing Oracle Inventory directory. Refer to "[Creating the Oracle Inventory Group](#)" on page 2-20 for more information.

Note: Oracle recommends that you do not put the `oraInventory` directory under Oracle base for a new installation. However, if you have an existing installation, then you should follow the steps suggested in this section.

- Identifying an existing Oracle home directory

Enter the following command to display the contents of the `oratab` file:

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

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

```
*:/u03/app/oracle/product/11.2.0/dbhome_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, then you can choose one of the following directories listed in the previous example:

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

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

- Identifying an existing Oracle base directory

After you have located the Oracle home directory, you can run the following command to confirm the location of Oracle base:

```
cat inventory/ContentsXML/oraclehomeproperties.xml
```

Before deciding to use an existing Oracle base directory for this installation, ensure 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 4 GB
The Oracle base directory will contain both software and database files (not recommended for production databases).	Up to 6 GB

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

```
# df -h oracle_base_path
```

To continue:

- If an Oracle base directory exists and you want to use it, then refer to the ["Choosing a Storage Option for Oracle Database and Recovery Files"](#) section on page 2-32.

When you configure 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 an Oracle base directory, then refer to the following section.

Creating an Oracle Base Directory

Before you create an Oracle base directory, you must identify an appropriate file system with sufficient free disk space.

To identify an appropriate file system:

1. To determine the free disk space on each mounted file system use the following command:

```
# df -h
```

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:

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
```

For example:

```
# mkdir -p /u01/app/oracle
# chown -R oracle:oinstall /u01/app/oracle
# chmod -R 775 /u01/app/oracle
```

2. When you configure the `oracle` user's environment later in this chapter, set the `ORACLE_BASE` environment variable to specify the Oracle base directory that you have created.

Choosing a Storage Option for Oracle Database and Recovery Files

Oracle Database files include data files, 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, then you must also choose the storage option that you want to use for recovery files (the fast recovery area). You do not have to use the same storage option for each file type.

Note: Database files and recovery files are supported on file systems and Automatic Storage Management.

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.
- Determine whether you want to use Automatic Storage Management for Oracle Database files, recovery files, or both. Refer to "[Step 1: Identifying Storage Requirements for Automatic Storage Management](#)" for more information.
- For more information about these storage options, refer to the "[Database Storage Options](#)" section on page 1-10.

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

- To use a file system for database or recovery file storage, refer to the "[Creating Directories for Oracle Database or Recovery Files](#)" section on page 2-33.
- To use Automatic Storage Management for database or recovery file storage, refer to the "[Preparing Disks for an Automatic Storage Management Installation](#)" section on page 3-5.

Creating Directories for Oracle Database or Recovery Files

This section contains the following topics:

- [Guidelines for Placing Oracle Database Files on a File System](#)
- [Creating Required Directories](#)

Guidelines for Placing Oracle Database Files on a File System

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

- The default path suggested by Oracle Universal Installer for the database file directory is a subdirectory of the Oracle base directory.
- 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, then 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, then choose file systems on separate physical devices that are dedicated to the database.

This method enables you to distribute physical input-output operations and create separate control files on different devices for increased reliability. It also enables you to fully implement the OFA guidelines. You can choose the Advanced database creation option to implement this method.

- If you intend to create a preconfigured database during the installation, then the file system (or file systems) that you choose must have at least 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.

Creating Required Directories

Note: You must perform 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:

1. Use the following to determine the free disk space on each mounted file system:

```
# df -h
```

2. 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 2 GB of free disk space ■ Two or more file systems with at least 2 GB of free disk space in total
Recovery files	Choose a file system with at least 2.4 GB of free disk space

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

3. Note the names of the mount point directories for the file systems that you identified.
4. 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
```

The default location for Database file directory is \$ORACLE_BASE/oradata.

- Recovery file directory (fast recovery area):

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

The default fast recovery area is \$ORACLE_BASE/recovery_area. However, Oracle recommends that you keep the fast recovery area on a separate physical disk than that of the database file directory. This will enable you use the fast recovery area to retrieve data if the disk containing oradata is unusable due to any reasons.

5. If you also want to use Automatic Storage Management for storage, then refer to the following section:

["Preparing Disks for an Automatic Storage Management Installation"](#) on page 3-5, else refer to the ["Stopping Existing Oracle Processes"](#) section on page 2-37.

Configuring Storage for Oracle Database Files Using Block Devices

This section describes how to configure Oracle Database files on block devices. Use the following procedure to create block device partitions:

1. Use `fdisk` to create disk partitions on block devices for database files.

If you intend to configure block devices and use Automatic Storage Management to manage data files, then create one partition for each disk comprising the whole disk, and go through the section *Configuring Disks for Automatic Storage Management with ASMLIB* in *Oracle Grid Infrastructure Installation Guide*.

2. Create or modify a rules file in `/etc/udev/rules.d`, to change the permissions of the datafiles from default `root` ownership.

Ensure that the file you create is appropriate for your distribution. For example, name this file `99-oracle.rules` on Asianux, Red Hat Enterprise Linux, and Oracle Enterprise Linux, or `40-oracle.rules` on SUSE Enterprise Server systems.

Example 2-1 Example of a Rules File With User oracle

```
/etc/udev/rules.d/99-oracle.rules
#
# ASM disks
KERNEL=="sdb[6-9]", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="sdb10", OWNER="oracle", GROUP="dba", MODE="0660"
```

Example 2-2 Example of a Rules File With User grid

```
/etc/udev/rules.d/99-oracle.rules
#
# ASM disks
KERNEL=="sdb[6-9]", OWNER="grid", GROUP="asmadmin", MODE="0660"
KERNEL=="sdb10", OWNER="grid", GROUP="asmadmin", MODE="0660"
```

See Also: Chapter 2, "Preparing Storage for ASM" in the *Oracle Database Storage Administrator's Guide* for information on preparing the storage subsystem before you configure Automatic Storage Management.

Configuring Disk Devices for Oracle Database

The `O_DIRECT` parameter enables direct read and writes to block devices, avoiding kernel overhead. With Oracle Database Release 10.2 and later, Oracle Database files are configured by default to use direct input/output.

With the 2.6 kernel or later for Red Hat Enterprise Linux, Oracle Enterprise Linux, and SUSE Enterprise Server, you must create a permissions file to maintain permissions on Oracle database files. If you do not create this permissions file, then permissions on disk devices revert to their default values, `root:disk`, and Oracle database fails to start. Use the following steps to set the permissions file number:

- On Asianux 2, Red Hat Enterprise Linux 4, and Oracle Enterprise Linux 4, you must create a permissions file number that is lower than 50.

- On Asianux 3, Red Hat Enterprise Linux 5, Oracle Enterprise Linux 5, or SUSE Enterprise Linux 10, you must create a permissions file number that is higher at 50.

To configure a permissions file for disk devices, complete the following tasks:

- [Example of Creating a Udev Permissions File for Oracle Database](#)
- [Example of Configuring Block Device Storage for Oracle Database](#)

See Also: *Oracle Grid Infrastructure Installation Guide* for information about configuring storage for Oracle database files on shared storage devices.

Example of Creating a Udev Permissions File for Oracle Database

Refer to the examples in "[Configuring Storage for Oracle Database Files Using Block Devices](#)" for more information on creating a permissions file.

Example of Configuring Block Device Storage for Oracle Database

The following is the procedure to create partitions for Oracle Database files on block devices:

1. Log in as root
2. Enter the `fdisk` command to format a specific storage disk. For example, `/sbin/fdisk /dev/sdb`
3. Create a partition. For instance, make a partition of 280 MB for data files.
4. Use the command similar to the following to update the kernel partition table for the shared storage device:

```
/sbin/partprobe diskpath
```

The following is an example of how to use `fdisk` to create one partition on a shared storage block disk device for a data file:

```
$ sudo sh
Password:
# /sbin/fdisk /dev/sdb
The number of cylinders for this disk is set to 1024.
Command (m for help): n
Command action
   e   extended
   P   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-1024, default 1):
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-4462, default 1)
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-1024, default 4462): using default
value 4462

Command (m for help):w

The partition table has been altered!
Calling ioctl () to re-read partition table.
Synching disks.
# exit
```

```
Last login Wed Feb 21 20:23:01 from localnode
$ sudo sh
Password:
# /sbin/partprobe /dev/sdb1
```

Stopping Existing Oracle Processes

Note:

- If you plan to use Oracle Restart, you must install the grid infrastructure before you install the database. As such, when you perform a database installation, then the database uses the same listener created during the grid infrastructure installation and you do not have to perform the steps listed in this section. Refer to "[Configuring the oracle User's Environment](#)" to continue.
 - If you are installing additional Oracle Database 11g products in an existing Oracle home, then stop *all* processes running in the Oracle home. You must complete this task to enable Oracle Universal Installer to relink certain executables and libraries.
-
-

If you choose to create a database during the installation, then most installation types configure and start a default Oracle Net listener using TCP/IP port 1521 and the IPC key value `EXTPROC`. However, if an existing Oracle Net listener process is using the same port or key value, Oracle Universal Installer will look for the next available free port (for example, 1522) and will configure and start the new listener on this available free port.

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

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, then refer to the "[Configuring the oracle User's Environment](#)" section on page 2-38 to continue.

3. Set the `ORACLE_HOME` environment variable to specify the appropriate Oracle home directory for the listener:
 - Bourne, Bash, or Korn shell:

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

- C or tcsh shell:

```
% setenv ORACLE_HOME oracle_home1
```

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

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

Note: If the listener uses the default name LISTENER, then 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.

Configuring the oracle User's Environment

You run Oracle Universal Installer from the `oracle` account. However, before you start Oracle Universal 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` environment variable.

To set the `oracle` user's environment:

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 fully_qualified_remote_host_name
```

For example:

```
$ xhost somehost.us.example.com
```

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


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


```
$ echo $SHELL
```
6. To run the shell startup script, enter one of the following commands:

- Bash shell:

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

- Bourne or Korn shell:

```
$ . ~/.profile
```

- C shell:

```
% source ~/.login
```

7. If you are not installing the software on the local computer, then run the following command on the remote machine to set the DISPLAY variable:

- Bourne, Bash or Korn shell:

```
$ export DISPLAY=local_host:0.0
```

- C shell:

```
% setenv DISPLAY local_host:0.0
```

In this example, `local_host` is the host name or IP address of the local computer that you want to use to display Oracle Universal Installer.

Run the following command on the remote machine to check if the shell and the DISPLAY environmental variable are set correctly:

```
echo $SHELL
echo $DISPLAY
```

Now to enable X applications, run the following commands on the local computer:

```
$ xhost + fully_qualified_remote_host_name
```

To verify that X applications display is set properly, run a X11 based program that comes with the operating system such as `xclock`:

```
$ xclock_path
```

In this example, `xclock_path` is the directory path. For example, you can find `xclock` at `/usr/X11R6/bin/xclocks`. If the DISPLAY variable is set properly, then you can see `xclock` on your computer screen.

See Also: PC-X Server or operating system vendor documents for further assistance

8. If you determined that the `/tmp` directory has less than 1 GB of free disk space, then identify a file system with at least 1 GB of free space and set the TMP and TMPDIR environment variables to specify a temporary directory on this file system:

- a. To determine the free disk space on each mounted file system use the following command:

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

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

```
$ sudo mkdir /mount_point/tmp
$ sudo chmod a+wr /mount_point/tmp
# exit
```

- c. Enter commands similar to the following to set the TMP and TMPDIR environment variables:

* Bourne, Bash, or Korn shell:

```
$ TMP=/mount_point/tmp
$ TMPDIR=/mount_point/tmp
$ export TMP TMPDIR
```

* C shell:

```
% setenv TMP /mount_point/tmp
% setenv TMPDIR /mount_point/tmp
```

9. 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).

10. Enter the following commands 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, then Oracle Universal 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, then Oracle recommends that you unset the ORACLE_HOME environment variable and choose the default path suggested by Oracle Universal Installer.

See Also: "[Configuring the User's Environment](#)" for information about setting the grid infrastructure software owner user's environment

Oracle Grid Infrastructure

The Oracle grid infrastructure for a standalone server is the Oracle software that enables a server to participate in an enterprise grid architecture. If you plan to use Oracle Restart or Automatic Storage Management, you must install the grid infrastructure before installing your database. The grid infrastructure for a standalone server is the software that includes Oracle Restart and Automatic Storage Management. Oracle combined the two infrastructure products into a single set of binaries that is installed as the grid infrastructure home. The grid infrastructure should be installed before installing Oracle Database 11g Release 2.

Automatic Storage Management is a volume manager and a file system for Oracle database files that supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations. Automatic Storage Management also supports a general purpose file system for your application needs including Oracle Database binaries. Automatic Storage Management is Oracle's recommended storage management solution that provides an alternative to conventional volume managers, file systems, and raw devices.

Oracle Restart improves the availability of your Oracle database by providing the following:

- When there is a hardware or a software failure, Oracle Restart automatically starts all Oracle components, including Oracle database instance, Oracle Net Listener, database services, and Automatic Storage Management.
- Oracle Restart starts up components in the proper order when the database host is restarted.
- Oracle Restart runs periodic checks to monitor the health of Oracle components. If a check operation fails for a component, then the component is shut down and restarted.

Note:

- If you want to use Automatic Storage Management or Oracle Restart, then you must first install Oracle grid infrastructure for a standalone server and then install Oracle Database.
 - Oracle Restart is used in single-instance (non-clustered) environments only.
-
-

This chapter contains the following sections:

- [Requirements for Oracle Grid Infrastructure Installation](#)
- [Oracle ACFS and Oracle ADVM Support](#)

- [Migrating Existing Automatic Storage Management Instances](#)
- [Automatic Storage Management Installation Considerations](#)
- [Preparing Disks for an Automatic Storage Management Installation](#)
- [Installing and Configuring Grid Infrastructure for a Standalone Server](#)
- [Manually Configuring Oracle Automatic Storage Management Disk Groups](#)
- [Testing the Automatic Storage Management Installation](#)

Requirements for Oracle Grid Infrastructure Installation

The system must meet the following requirements:

- [Memory Requirements](#)
- [Disk Space Requirements](#)
- [Configuring the User's Environment](#)

Memory Requirements

The following are the memory requirements for installing Oracle grid infrastructure:

- At least 1 GB of RAM

To determine the RAM size, enter the following command:

```
# grep MemTotal /proc/meminfo
```

If the size of the RAM is less than the required size, then you must install more memory before continuing.

- The following table describes the relationship between installed RAM and the configured swap space requirement:

Note: On Linux, the HugePages feature allocates non-swappable memory for large page tables using memory-mapped files. If you enable HugePages, then you should deduct the memory allocated to HugePages from the available RAM before calculating swap space.

RAM	Swap Space
Between 1 GB and 2 GB	1.5 times the size of RAM
Between 2 GB and 16 GB	Equal to the size of RAM
More than 16 GB	16 GB

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

```
# grep SwapTotal /proc/meminfo
```

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

To determine the available RAM and swap space, enter the following command:

```
# free
```

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

Disk Space Requirements

The following are the disk space requirements for installing grid infrastructure:

- At least 2.2 GB of disk space.
- At least 1 GB of disk space in the `/tmp` directory.

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

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

If there is less than 1 GB of free disk space available in the `/tmp` directory, then complete one of the following steps:

- Delete unnecessary files from the `/tmp` directory to meet the disk space requirement.
- Set the `TMP` and `TMPDIR` environment variables when setting the `oracle` user's environment.

See Also: ["Configuring the oracle User's Environment"](#) for more information on setting `TMP` and `TMPDIR`

- Extend the file system that contains the `/tmp` directory. If necessary, contact the system administrator for information about extending file systems.

Configuring the User's Environment

Complete the following tasks to set the Oracle grid infrastructure software owner user's environment:

- Review the information in ["Logging In to the System as root"](#) section.
- Ensure that you set the path to the Oracle base directory. Oracle Restart and Oracle Database reside under the same Oracle base. For example:

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

- Set the Oracle grid infrastructure software owner user default file mode creation mask (`umask`) to `022` in the shell startup file. Setting the mask to `022` ensures that the user performing the software installation creates files with `755` permissions.
- Set `ulimit` settings for file descriptors and processes for the Oracle grid infrastructure installation software owner.
- Set the `DISPLAY` environment variable in preparation for installation.

Oracle ACFS and Oracle ADVM Support

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

Management Dynamic Volume Manager (Oracle ADVN) provides volume management services and a standard disk device driver interface to clients. Oracle Automatic Storage Management Cluster File System is layered on Automatic Storage Management through the Oracle Automatic Storage Management Dynamic Volume Manager interface.

Oracle Automatic Storage Management Cluster File System and Oracle Automatic Storage Management Dynamic Volume Manager are supported only on Oracle Enterprise Linux 5.0 and Red Hat Enterprise Linux 5.0 for Linux x86 and Linux x86-64.

Note: Oracle recommends that Oracle data files are installed in Oracle ASM disk groups. Installing Oracle data files on an Oracle ACFS file system is not supported. Oracle ACFS can be used as an option only when Automatic Storage Management is configured.

See Also:

- *Oracle Database Release Notes for Linux* for latest information on supported platforms and releases
- *Oracle Database Storage Administrator's Guide* for more information about Oracle Automatic Storage Management Cluster File System and Oracle Automatic Storage Management Dynamic Volume Manager

Migrating Existing Automatic Storage Management Instances

If you have an Automatic Storage Management (ASM) installation from a prior release installed on your server, or in an existing Oracle grid infrastructure installation, you can use Oracle Automatic Storage Management Configuration Assistant (ASMCA) to upgrade the existing Automatic Storage Management instance to 11g Release 2 (11.2), and subsequently configure disk groups, Automatic Storage Management volumes and Automatic Storage Management file systems.

Note: You must first shut down all databases and applications using an existing Automatic Storage Management instance before upgrading it.

During installation, if you chose to use Automatic Storage Management and Oracle Automatic Storage Management Configuration Assistant detects that there is a prior Automatic Storage Management version installed in another Automatic Storage Management home, then after installing the Automatic Storage Management 11g Release 2 (11.2) binaries, you can start Oracle Automatic Storage Management Configuration Assistant to upgrade the existing Automatic Storage Management instance. You can then choose to configure an Automatic Storage Management File System deployment by using the Automatic Storage Management Configuration Assistant to create an Oracle Automatic Storage Management Dynamic Volume and an Oracle Automatic Storage Management Cluster File System. (Oracle ACFS).

See Also: ■

- "Upgrading an Oracle ASM Instance with Oracle ASM Configuration Assistant" in *Oracle Database Storage Administrator's Guide*
- "Upgrading an Oracle ASM Instance With Oracle Universal Installer" in *Oracle Database Storage Administrator's Guide*

Automatic Storage Management Installation Considerations

In previous releases, Oracle Automatic Storage Management (ASM) was installed as part of the Oracle Database installation. With Oracle Database 11g Release 2 (11.2), Automatic Storage Management is part of an Oracle grid infrastructure installation, either for a cluster, or for a standalone server.

If you want to upgrade an existing Automatic Storage Management installation, then you must upgrade Automatic Storage Management by running an Oracle grid infrastructure upgrade. If you do not have Automatic Storage Management installed and you want to use Automatic Storage Management as your storage option, then you must complete an Oracle grid infrastructure installation before you start your Oracle Database installation

Apply the following guidelines when you install Automatic Storage Management:

- You must complete the steps listed under "[Preparing Disks for an Automatic Storage Management Installation](#)" to prepare a disk partition to use for the Automatic Storage Management disk groups.
- Ensure that at least one disk group exists before beginning the installation.
- When you install Automatic Storage Management, Oracle Automatic Storage Management Configuration Assistant creates a separate server parameter file (*SPFILE*) and password file for the Automatic Storage Management instance.

Preparing Disks for an Automatic Storage Management Installation

This section describes how to configure disks for use with Automatic Storage Management. 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:

- [General Steps for Configuring Automatic Storage Management](#)
- [Step 1: Identifying Storage Requirements for Automatic Storage Management](#)
- [Step 2: Using an Existing Automatic Storage Management Disk Group](#)
- [Step 3: Creating DAS or SAN Disk Partitions for Automatic Storage Management](#)
- [Step 4: Configuring Disks for Automatic Storage Management](#)

Note: Oracle does not recommend using identifiers for database object names that must be quoted. While these quoted identifiers may be valid as names in the SQL `CREATE` statement, such as `CREATE DISKGROUP "1data" . . .`, the names may not be valid when using other tools that manage the database object.

See Also: "Creating Disk Groups for a New Oracle Installation", in *Oracle Database Storage Administrator's Guide* for information on creating and managing disk groups

General Steps for Configuring Automatic Storage Management

The following are the general steps to configure Automatic Storage Management:

1. Identify the storage requirements of the site.
2. Optionally, use an existing Automatic Storage Management disk group.
3. If you are creating a new Automatic Storage Management disk group, create partitions for DAS or SAN disks.
4. Use one of the following methods to complete the Automatic Storage Management configuration:
 - If you plan to install Oracle Database using interactive mode, Oracle Universal Installer prompts you for the Automatic Storage Management disk configuration information during the installation.
 - If you plan to install Oracle Database using silent or response file mode, you will need to manually configure the disks before performing the installation.

Step 1: Identifying Storage Requirements for Automatic Storage Management

To identify the storage requirements for using Automatic Storage Management, you must determine the number of devices and the amount of free disk space that you require. To complete this task:

1. Determine whether you want to use Automatic Storage Management for Oracle Database files, recovery files, or both.

Note: You do not have to use the same storage mechanism for , Oracle Database files, and recovery files. You can use a file system for one file type and Automatic Storage Management for the other.

If you choose to enable automated backups and you do not have a shared file system available, then you must choose Automatic Storage Management for recovery file storage.

If you plan to enable automated backups during the installation, then you can choose Automatic Storage Management as the storage mechanism for recovery files by specifying an Automatic Storage Management disk group for the fast 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 Oracle Automatic Storage Management Configuration Assistant in interactive mode, by choosing the Advanced database configuration option for example, then you can decide whether you want to use the same Automatic Storage Management disk group for database files and recovery files, or you can choose to use different disk groups for each file type. Ideally, you should create separate Automatic Storage Management disk groups for data files and for recovery files.

The same choice is available to you if you use Oracle Database Configuration Assistant after the installation to create a database.

- If you select an installation type that runs Oracle Automatic Storage Management Configuration Assistant in silent or response file mode, then you must use the same Automatic Storage Management disk group for data files and recovery files.

See Also: ["Create a Fast Recovery Area Disk Group"](#)

2. Choose the Automatic Storage Management redundancy level that you want to use for each Automatic Storage Management disk group that you create.

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

- External redundancy

This option does not allow Automatic Storage Management to mirror the contents of the disk group. Oracle recommends that you select this redundancy level either when the disk group contains devices, such as RAID devices, that provide their own data protection or when the database does not require an uninterrupted access to data.

- Normal redundancy

To optimize performance and reliability in a normal redundancy disk group, Automatic Storage Management uses two-way mirroring for data files and three-way mirroring for control files, by default. In addition, you can choose the mirroring characteristics for individual files in a disk group. Alternatively, you can use two-way mirroring or no mirroring.

A normal redundancy disk group requires a minimum of two failure groups (or two disk devices) if you are using two-way mirroring. 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

The contents of the disk group are three-way mirrored by default. To create a disk group with high redundancy, you must specify at least three failure groups (a minimum of 3 devices).

Although high-redundancy disk groups 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.

If an Automatic Storage Management instance is already running on the system, then you can use an existing disk group to meet these storage requirements. If necessary, you can add disks to an existing disk group during the installation.

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	Data Files	Recovery Files	Both File Types
External	1	1.8 GB	3.6 GB	5.4 GB
Normal	2	3.6 GB	7.2 GB	10.8 GB
High	3	5.4 GB	10.8 GB	16.2 GB

- Optionally, identify failure groups for the Automatic Storage Management disk group devices.

If you intend to use a normal or high redundancy disk group, then you can further protect the database against hardware failure by associating a set of disk devices in a custom failure group. By default, each device comprises its failure group. However, if two disk devices in a normal redundancy disk group are attached to the same SCSI controller, then the disk group becomes unavailable if the controller fails. The controller in this example is a single point of failure.

To avoid failures of this type, you can 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, then you must specify a minimum of two failure groups for normal redundancy disk groups and three failure groups for high redundancy disk groups.

- If you are sure that a suitable disk group does not exist on the system, then install or identify appropriate disk devices to add to a new disk group. Apply the following guidelines when identifying appropriate disk devices:

- The disk devices must be owned by the user performing the grid installation.

See Also: [Example 2-2, "Example of a Rules File With User grid"](#) for information on creating or modifying permissions

- All the devices in an Automatic Storage Management disk group should be the same size and have the same performance characteristics.
- Do not specify multiple partitions on a single physical disk as a disk group device. Automatic Storage Management expects each disk group device to be on a separate physical disk.
- Oracle does not recommend the use of a logical volume as a device in Automatic Storage Management because the logical volume is capable of hiding the physical disk architecture which prevents Automatic Storage Management from optimizing I/O across physical devices.

See Also: ■

- "[Step 4: Configuring Disks for Automatic Storage Management](#)" on page 3-10 for information about completing this task
- "Preparing Storage for ASM" in *Oracle Database Storage Administrator's Guide* for information on configuring Automatic Storage Management disk groups

Step 2: Using an Existing Automatic Storage Management Disk Group

Note: This is an optional step.

This step describes how to identify existing disk groups and determine the free disk space that they contain. If you want to store either database or recovery files in an existing Automatic Storage Management disk group, then you have the following choices, depending on the installation method that you select:

- Before you run Oracle Automatic Storage Management Configuration Assistant in interactive mode, ensure that you have already created a disk group or selected an existing one.
- If you select an installation method that runs Oracle Automatic Storage Management Configuration Assistant in silent or response file mode, then you must choose an existing disk group for the new database; you cannot create a disk group. However, you can add disk devices to an existing disk group if it has insufficient free space for your requirements.

Note: The Automatic Storage Management instance that manages the existing disk group will be running in the grid infrastructure Oracle home directory.

To determine whether an existing Automatic Storage Management 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 Automatic Storage Management instance is configured on the system:

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

If an Automatic Storage Management instance is configured on the system, then 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 Oracle Automatic Storage Management instance and `oracle_home_path` is the Oracle home directory where Oracle Automatic Storage Management is installed. By convention, the `SID` for an Oracle Automatic Storage Management instance should be `+ASM`.

2. Open a shell prompt and temporarily set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the Automatic Storage Management instance that you want to use.

For example, the Automatic Storage Management `SID`, which is named `+ASM` and is located in the `grid` subdirectory of the `ORACLE_BASE` directory, then enter the following commands to create the required settings:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_SID=+ASM
$ export ORACLE_SID
$ ORACLE_HOME=/u01/app/oracle/product/11.2.0/grid/
$ export ORACLE_HOME
```

- C shell:

```
% setenv ORACLE_SID +ASM
% setenv ORACLE_HOME /u01/app/oracle/product/11.2.0/grid
```

3. By using SQL*Plus, connect to the Automatic Storage Management instance as the SYS user with SYSASM privilege and start the instance if necessary:

```
# $ORACLE_HOME/bin/sqlplus
SQL> CONNECT SYS as SYSASM
Enter password: SYS_password
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, then Oracle recommends that you use devices that have the same size and performance characteristics as the existing devices in that disk group.

Step 3: Creating DAS or SAN Disk Partitions for Automatic Storage Management

In order to use a DAS or SAN disk in Automatic Storage Management, the disk must have a partition table. Oracle recommends creating exactly one partition for each disk containing the entire disk.

Note: You can use any physical disk for Automatic Storage Management, as long as it is partitioned.

Step 4: Configuring Disks for Automatic Storage Management

Oracle provides an Automatic Storage Management library driver that you can use to simplify the configuration and management of the disk devices that you want to use with Automatic Storage Management. A disk that is configured for Automatic Storage Management is known as a candidate disk.

If you intend to use Automatic Storage Management for database storage on Linux, then Oracle recommends that you install the Automatic Storage Management library driver (ASMLIB) and associated utilities and use them to configure the devices that you want to include in an Automatic Storage Management disk group.

Note: If you choose to configure disks using the Automatic Storage Management library driver, then you must change the default disk discovery string to ORCL: *. If the diskstring is set to ORCL: *, or is left empty (""), then the installer discovers these disks.

Configuring Disks for Automatic Storage Management Using the Automatic Storage Management Library Driver (ASMLIB)

To use the Automatic Storage Management library driver to configure Automatic Storage Management devices, complete the following tasks:

- [Installing and Configuring the Automatic Storage Management Library Driver Software](#)
- [Configuring the Disk Devices to Use the Automatic Storage Management Library Driver](#)
- [Administering the Automatic Storage Management Library Driver and Disks](#)

Installing and Configuring the Automatic Storage Management Library Driver Software

To install and configure the Automatic Storage Management library driver software:

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

```
# uname -rm
```

2. If necessary, download the required Automatic Storage Management library driver packages from the Oracle Technology Network Web site:

<http://www.oracle.com/technology/tech/linux/asmlib/index.html>

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

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

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

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

For example, if you are using the Red Hat Enterprise Linux AS 3.0 enterprise kernel on an x86 system, then enter a command similar to the following:

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

4. Enter a command similar to the following to determine the UID of the Oracle software owner user that you are using for this installation and the GID of the OSASM group:

```
# id oracle
```

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

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

6. 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 grid infrastructure owner user (typically, <code>grid</code>).
Default GID to own the driver interface:	Specify the GID of the OSASM group (typically, <code>asmadmin</code>).
Start Oracle Automatic Storage Management Library driver on start (y/n):	Enter <code>y</code> to start the Oracle Automatic Storage Management library driver when the system starts.
Scan for Oracle ASM disks on boot (y/n):	Enter <code>y</code> to scan for presence of any Oracle Automatic Storage Management disks when the system starts.

Configuring the Disk Devices to Use the Automatic Storage Management Library Driver

To configure the disk devices that you want to use in an Automatic Storage Management disk group:

1. If you intend to use IDE, SCSI, or RAID devices in the Automatic Storage Management disk group, then:
 - a. If necessary, install or configure the disk devices that you intend to use for the disk group and restart the system.
 - b. To identify the device name for the disks that you want to use, enter the following command:

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

Depending on the type of disk, the device name can vary:

Disk Type	Device Name Format	Description
IDE disk	<code>/dev/hdxn</code>	In this example, <code>x</code> is a letter that identifies the IDE disk and <code>n</code> is the partition number. For example, <code>/dev/hda</code> is the first disk on the first IDE bus.
SCSI disk	<code>/dev/sdxn</code>	In this example, <code>x</code> is a letter that identifies the SCSI disk and <code>n</code> is the partition number. For example, <code>/dev/sda</code> is the first disk on the first SCSI bus.
RAID disk	<code>/dev/rd/cxdypz</code> <code>/dev/ida/cxdypz</code>	Depending on the RAID controller, RAID devices can have different device names. In the examples shown, <code>x</code> is a number that identifies the controller, <code>y</code> is a number that identifies the disk, and <code>z</code> is a number that identifies the partition. For example, <code>/dev/ida/c0d1</code> is the second logical drive on the first controller.

Note: 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. Enter a command similar to the following to mark a disk as an Automatic Storage Management disk:

```
# /etc/init.d/oracleasm createdisk DISK1 /dev/sdb1
```

In this example, DISK1 is a name that you want to assign to the disk.

Note:

- If you are using a multipathing disk driver with Automatic Storage Management, then ensure that you specify the correct logical device name for the disk.

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

- To create a database during the installation using the Automatic Storage Management library driver, you must change the default disk discovery string to ORCL: *.
-
-

Administering the Automatic Storage Management Library Driver and Disks

To administer the Automatic Storage Management 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 Automatic Storage Management 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 Automatic Storage Management library driver when the system starts. The <code>enable</code> option causes the Automatic Storage Management library driver to load when the system starts: # /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 Automatic Storage Management library driver without restarting the system: # /etc/init.d/oracleasm restart
createdisk	Use the <code>createdisk</code> option to mark a disk device for use with the Automatic Storage Management library driver and give it a name: # /etc/init.d/oracleasm createdisk <i>DISKNAME devicename</i>
deletedisk	Use the <code>deletedisk</code> option to unmark a named disk device: # /etc/init.d/oracleasm deletedisk <i>DISKNAME</i>
	Note: Do not use this command to unmark disks that are being used by an Automatic Storage Management disk group. You must drop the disk from the Automatic Storage Management disk group before you unmark it.

Option	Description
querydisk	Use the <code>querydisk</code> option to determine whether a disk device or disk name is being used by the Automatic Storage Management library driver: # /etc/init.d/oracleasm querydisk {DISKNAME devicename}
listdisks	Use the <code>listdisks</code> option to list the disk names of marked Automatic Storage Management 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 Automatic Storage Management library driver disks on another node: # /etc/init.d/oracleasm scandisks

Installing and Configuring Grid Infrastructure for a Standalone Server

If you install Oracle grid infrastructure and then create your database, the database is automatically added to the Oracle grid infrastructure configuration, and is then automatically restarted when required. However, if you install Oracle grid infrastructure on a host computer on which a database already exists, you must manually add the database, the listener, the Oracle Automatic Storage Management instance, and other components to the Oracle grid infrastructure configuration.

Note: Oracle grid infrastructure can accommodate multiple single-instance databases on a single host computer.

This section includes the following topics:

- [Installing Grid Infrastructure with a New Database Installation](#)
- [Installing Grid Infrastructure for an Existing Database](#)

Installing Grid Infrastructure with a New Database Installation

Perform the following steps to install the grid infrastructure and then create a database that is managed by Oracle Restart. First install grid infrastructure, which installs Oracle Restart and Automatic Storage Management, then configure Automatic Storage Management with at least one disk group, and then install Oracle database that stores database files in Automatic Storage Management disk groups. Click the help button on the Oracle Universal Installer page for page level assistance.

To install Oracle grid infrastructure for a standalone server with a new database installation:

1. Start Oracle Universal Installer as the grid infrastructure software owner user. Complete one of the following steps depending on the location of the installation files:
 - If the installation files are on installation media, enter commands similar to the following, where `directory_path` is the path of the grid infrastructure directory on the installation media:

```
$ /directory_path/runInstaller
```

Note: You must install Oracle grid infrastructure for a standalone server from the grid infrastructure media.

- If the installation files are on the hard disk, change the directory to the path of the grid infrastructure (`clusterware`) directory and enter the following command:

```
$ ./runInstaller
```

Note: Start Oracle Universal Installer from the terminal session where you logged in as the grid infrastructure software owner user and set the user's environment.

See Also: ["Configuring the User's Environment"](#) for information about setting the grid infrastructure software owner user's environment

If Oracle Universal Installer is not displayed, refer to ["X Window Display Errors"](#) and ["Remote Terminal Installation Error"](#) for information about troubleshooting.

2. The Select Installation Option screen appears. Select **Install and Configure Grid Infrastructure for a Standalone Server** option, to install and configure Oracle Restart and Oracle Automatic Storage Management. Click **Next**.
3. On the Select Product Languages screen, select one or more languages. Move the languages from the Available Languages list to the Selected Languages list. Click **Next**.
4. The Create ASM Disk Group screen, lists all the Automatic Storage Management disks under `ORCL: *`

Click **Change Disk Discovery Path** to select any devices that will be used by Automatic Storage Management but are not listed. In the Change Disk Discovery Path window, enter a string to use to search for devices that Automatic Storage Management will use. If the diskstring is set to `ORCL: *` or is left empty (""), then the installer discovers these disks. Click **OK**.

After you finish selecting the disks to be used by Automatic Storage Management, click **Next**.

Note: During installation, disk paths mounted on ASM and registered on ASMLIB with the string `ORCL: *` are listed as default database storage candidate disks.

Consider the following information about disk devices while performing this step:

- The Disk Group Name default is `DATA`. You can enter a new name for the disk group, or use the default name.
- The disk devices must be owned by the user performing the grid installation.

See Also: [Example 2-2, "Example of a Rules File With User grid"](#) for information on creating or modifying permissions

- Check with your system administrator to determine if the disks used by Automatic Storage Management are mirrored at the storage level. If so, select External for the redundancy. If the disks are not mirrored at the storage level, then choose Normal for the redundancy.

Note: For normal redundancy, you require twice as much disk space to hold the same amount of data. For example, if your database is 100 GB, then you require approximately 200 GB of storage.

5. In the Specify ASM Password screen, enter SYSASM password required to connect to the Automatic Storage Management instance. The Oracle Automatic Storage Management instance is managed by a privileged role called SYSASM, which grants full access to Automatic Storage Management disk groups. Oracle recommends that you create a less privileged user, ASMSNMP, with SYSDBA privileges to monitor the Oracle Automatic Storage Management instance.

Enter passwords for the SYS and ASMSNMP user accounts. The passwords should be at least eight characters in length and include at least one alphabetic and one numeric character.

6. On the Privileged Operating System Groups screen, select the name of the operating system group you created for the OSDBA group, the OSASM group, and the database operator group OSOPER. If you choose to create only the dba group, then you can use that group for all three privileged groups. If you created a separate asmadmin group, then use that value for the OSASM group. Click **Next**.
7. On the Specify Installation Location screen, enter the following details and click **Next**:
 - Oracle Base: Enter the directory location for Oracle base. Do not include spaces in the path name.
 - Software Location: This field is populated by default in concurrence with Oracle base location.

See Also: ["Naming Directories"](#) for directory naming conventions

8. If you have not installed any Oracle software previously on this server, the Create Inventory screen appears.

Change the path for the Inventory Directory, if required. Select oinstall for the oraInventory Group Name, if required. Click **Next**.

9. The Perform Prerequisite Checks screen, checks if the minimum system requirements are met to carry out the grid infrastructure installation. If all the system requirements are met, then you will be directed to the Summary screen. However, in case of a failed installation, you can review the error.

If you click **Check Again**, then you can run the prerequisite check again to see if the minimum requirements are met to carry on with the database installation.

Click **Fix & Check Again**, if you want the installer to fix the problem and check the system requirements once more.

Note: The Fix & Check Again option generates a script that you need to run as the `root` user. This generated script sets some of the system parameters to Oracle-recommended values. Oracle recommends that you do not modify the contents of this script. Refer to "[Installation Fixup Scripts](#)" for more information on fixup scripts.

To get a list of failed requirements, select **ShowFailed** from the list. To get a list of all the prerequisite checks run by the OUI, select **Show All**. To get a list of the prerequisite checks that are successful, select **Show Succeeded**.

Note: To ignore all the errors and proceed with the database installation, check **Ignore All**.

10. Review the contents of the Summary screen, and click **Finish**.
11. The Setup screen displays the progress of the grid infrastructure installation. During the installation process, the Execute Configuration Scripts window appears. Do *not* click **OK** until you have run the scripts mentioned in this screen.

Run the `root.sh` and, if required, the `oraInstRoot.sh` configuration scripts as the `root` user.
12. The Finish screen displays the installation status. Click **Close** to end the installation, then **Yes** to confirm that you want to exit Oracle Universal Installer.

If you encounter any problems, refer to the configuration log for information. The path to the configuration log is displayed on the Configuration Assistants window.
13. If you want to create additional disk groups, then run the Oracle Automatic Storage Management Configuration Assistant (ASMCA) utility. For example, you can create another disk group named `RECOVERY` to store the fast recovery area.

Note: To check if the Oracle High Availability Service is installed properly, run `./crsctl check has` command from `Grid_home/bin` directory.

`Grid_home` is the path to the Oracle grid infrastructure home for a standalone server. `ohasd` is a daemon installed with Oracle grid infrastructure that starts software services, such as Oracle Automatic Storage Management.

14. Install Oracle database.

Because you have already installed Oracle grid infrastructure, Oracle Universal Installer will automatically add the database instance to the Oracle grid infrastructure configuration if you choose to create a database as part of the installation.

See Also: "[Installing the Oracle Database Software](#)" for information about installing Oracle database

Note: If a new database is installed after a grid infrastructure installation, then the listener runs from the grid infrastructure home. Because Oracle Automatic Storage Management is installed as part of the grid infrastructure, the default listener is created and runs from the grid infrastructure home. If you perform a database installation, then the database uses the same listener created during the grid infrastructure installation.

Installing Grid Infrastructure for an Existing Database

Follow the high-level instructions in this section to install Oracle grid infrastructure and configure it for an existing Oracle database. To install Oracle grid infrastructure for an existing database:

- On the same host computer as the database, use Oracle Universal Installer to install the grid infrastructure, and select **Install and Configure Grid Infrastructure** for a Standalone Server as the installation option.

The grid infrastructure components are installed in a separate Oracle home.

Refer to ["Installing Grid Infrastructure with a New Database Installation"](#) on page 14 for detailed instructions.

- Go to the Grid home `bin` directory. Use the `srvctl add database` command to manually add the database, the listener, the Automatic Storage Management instance, all Automatic Storage Management disk groups, and any database services to the Oracle grid infrastructure configuration.

See Also: "srvctl add database" in *Oracle Database Administrator's Guide* for more information about the `srvctl add database` command

Manually Configuring Oracle Automatic Storage Management Disk Groups

The Oracle Automatic Storage Management Configuration Assistant utility creates a new Automatic Storage Management instance if there is no Automatic Storage Management instance currently configured on this computer. After installing the grid infrastructure, you can also use Oracle Automatic Storage Management Configuration Assistant to create and configure disk groups, Oracle Automatic Storage Management Dynamic Volume Manager. (Oracle ADVM) and Oracle Automatic Storage Management Cluster File System (Oracle ACFS). If you want to create additional disk groups or manually configure Automatic Storage Management disks, then you can run the Oracle Automatic Storage Management Configuration Assistant as follows:

```
$ cd Grid_home/bin
$ ./asmca
```

Where *Grid_home* is the path to the Oracle grid infrastructure home for a standalone server.

See Also: *Oracle Database Storage Administrator's Guide* for further information on Oracle Automatic Storage Management Configuration Assistant

Testing the Automatic Storage Management Installation

To test the Automatic Storage Management installation, try logging in by using the `asmcmd` command-line utility, which lets you manage Automatic Storage Management disk group files and directories. To do this:

1. Open a shell window, and temporarily set the `ORACLE_SID` and `ORACLE_HOME` environment variables to specify the appropriate values for the Automatic Storage Management instance that you want to use.

For example, if the Automatic Storage Management `SID` is named `+ASM` and the Oracle home is located in the `grid` subdirectory of the `ORACLE_BASE` directory, then enter the following commands to create the required settings:

- Bourne, Bash, or Korn shell:

```
$ ORACLE_SID=+ASM
$ export ORACLE_SID
$ ORACLE_HOME=/u01/app/oracle/product/11.2.0/grid
$ export ORACLE_HOME
```

- C shell:

```
% setenv ORACLE_SID +ASM
% setenv ORACLE_HOME /u01/app/oracle/product/11.2.0/grid
```

2. Use `asmcmd` to connect to the Automatic Storage Management instance as the `SYS` user with `SYSASM` privilege and start the instance if necessary:

```
$ORACLE_HOME/bin/asmcmd lsdg
```

See Also:

- *Oracle Database Utilities* for more information about `asmcmd`
- *Oracle Database Storage Administrator's Guide* for a more detailed description of Automatic Storage Management

Installing Oracle Database

The Oracle Database software is available on installation media or you can download it from the Oracle Technology Network Web site. In most cases, you use the graphical user interface (GUI) provided by Oracle Universal Installer to install the software. However, you can also use Oracle Universal Installer to complete silent-mode installations, without using the GUI.

- [Preinstallation Considerations](#)
- [Reviewing Component-Specific Installation Guidelines](#)
- [Accessing the Installation Software](#)
- [Database Security Options](#)
- [Installing the Oracle Database Software](#)
- [Installing Oracle Database Examples](#)

See Also: [Appendix A](#) for information about silent-mode installations

Preinstallation Considerations

After reviewing the information in [Chapter 1, "Overview of Oracle Database Installation"](#) and completing the tasks listed in [Chapter 2, "Oracle Database Preinstallation Requirements"](#), consider the following case.

Performing Multiple Oracle Database Installations in Response File or Silent Mode

If you must perform multiple installations of Oracle Database, you may want to use silent or response file mode. In response file mode, at each node, you run Oracle Universal Installer from the command line using a response file. The response file is a text file containing the settings you normally enter in the Oracle Universal Installer GUI dialog boxes.

See Also: [Appendix A](#) for information about silent mode installations

Reviewing Component-Specific Installation Guidelines

Review the following guidelines before starting Oracle Universal Installer:

- Oracle Universal Installer
Using Oracle Universal Installer from an earlier Oracle release to install components from this release is no longer allowed.

- Automatic Storage Management

In previous releases, Oracle Automatic Storage Management (ASM) was installed as part of the Oracle Database installation. With Oracle Database 11g Release 2 (11.2), Automatic Storage Management is part of an Oracle grid infrastructure installation, either for a cluster, or for a standalone server.

If you want to upgrade an existing Automatic Storage Management installation, then you must upgrade Automatic Storage Management by running an Oracle grid infrastructure upgrade. If you do not have Automatic Storage Management installed and you want to use Automatic Storage Management as your storage option, then you must complete an Oracle grid infrastructure installation before you start your Oracle Database installation.

See Also: ["Oracle Grid Infrastructure"](#) for information about Oracle grid infrastructure for a standalone server

- Installations on a Cluster

If Oracle Clusterware and Oracle RAC are already installed on the system, Oracle Universal Installer displays the Specify Hardware Cluster Installation page. You must select the Local Installation option, unless you want to install Oracle RAC.

See Also: *Oracle Real Application Clusters Installation Guide* for information about installing Oracle RAC

Selecting the Database Character Set

Oracle Database uses the database character set for:

- Data stored in SQL character datatypes (CHAR, VARCHAR2, CLOB, and LONG).
- Identifiers such as table names, column names, and PL/SQL variables.
- Stored SQL and PL/SQL source code, including text literals embedded in this code.

Once a database is created, changing its character set is usually very expensive in terms of time and resources. Such operation may require converting all character data by exporting the whole database and importing it back. Therefore, it is important that you carefully select the database character set already at installation time.

Oracle recommends Unicode AL32UTF8 as the database character set. Unicode is the universal character set that supports most of the currently spoken languages of the world. It also supports many historical scripts (alphabets). Unicode is the native encoding of many technologies, including Java, XML, XHTML, ECMAScript, and LDAP. Unicode is ideally suited for databases supporting the Internet and the global economy.

As AL32UTF8 is a multibyte character set, database operations on character data may be slightly slower when compared to single-byte database character sets, such as WE8MSWIN1252. Storage space requirements for text in most languages that use characters outside of the ASCII repertoire are higher in AL32UTF8 compared to legacy character sets supporting the language. Note that the increase in storage space concerns only character data and only data that is not in English. The universality and flexibility of Unicode usually outweighs these additional costs.

Legacy character sets should be considered when compatibility, storage requirements, or performance of text processing is critical and the database will ever support only a single group of languages. The database character set to be selected in such case is the character set of most clients connecting to this database.

The default character set suggested or used by Oracle Universal Installer and Database Configuration Assistant in this release is based on the language configuration of the operating system.

For most languages, the default character set is one of the Microsoft Windows character sets, for example WE8MSWIN1252, even though the database is not installed on Windows. This results from the assumption that most clients connecting to the database run under the Microsoft Windows operating system. As the database should be able to store all characters coming from the clients and Microsoft Windows character sets have richer character repertoire than the corresponding ISO 8859 character sets, the Microsoft Windows character sets are usually the better choice. For example, the EE8MSWIN1250 character set supports the Euro currency symbol and various smart quote characters, while the corresponding EE8ISO8859P2 character set does not support them. In any case, Oracle converts the data between the database character set and the client character sets, which are declared by the NLS_LANG settings.

The list of database character sets that is presented to you for selection by Oracle Universal Installer contains only the recommended character sets. Even though Oracle Database supports many more character sets, they are either deprecated or they are binary subsets of another recommended character set. For example, WE8DEC is a deprecated character set and US7ASCII and WE8ISO8859P1 are both binary subsets of WE8MSWIN1252.

If, for compatibility reasons, you must create a database in one of the non-recommended character sets, choose the Advanced database configuration option. Database Configuration Assistant in the interactive mode will give you the opportunity to select any of the database character sets supported on Linux.

Accessing the Installation Software

The Oracle Database software is available on installation media or you can download it from the Oracle Technology Network Web site, or Oracle E-Delivery Web site. To install the software from the hard disk, you must either download it and unpack it, or copy it from the installation media, if you have it.

You can access and install Oracle Database by using one of the following methods:

- To copy the software to a hard disk, refer to ["Copying the Software to the Hard Disk"](#) on page 4-5
- To download the software from Oracle Technology Network, refer to ["Downloading Oracle Software"](#) on page 4-3

Downloading Oracle Software

You can download the trial version of the installation files from the Oracle Technology Network (OTN) or Oracle E-Delivery web site and extract them on your hard disk. Make sure that you completely review and understand the terms of the license. Most downloads include the Development License. This section contains the following topics:

- [Downloading the Installation Archive Files from OTN](#)
- [Downloading the Software from Oracle E-Delivery](#)
- [Extracting the Installation Files](#)

Downloading the Installation Archive Files from OTN

To download the installation archive files from Oracle Technology Network:

1. Use any browser to access the software download page from Oracle Technology Network:
<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 file names.
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 `OraDB11g`, 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 Oracle Technology Network.
8. Extract the files in each directory that you just created.
9. After you have extracted the required installation files, refer to "[Installing the Oracle Database Software](#)" on page 4-8.

Downloading the Software from Oracle E-Delivery

You can download the software from Oracle E-Delivery as Media Packs. A Media Pack is an electronic version of the software that is also available to Oracle customers on CD-ROM or DVD. To download the Media Pack:

1. Use any browser to access the Oracle E-Delivery Web site:
<http://edelivery.oracle.com/>
2. Complete the Export Validation process by entering a minimal amount of information (name, company, e-mail address, and country) in the online form.
3. In the Media Pack Search page, specify the Product Pack and Platform to identify the Media Pack you want to download. If you do not know the name of the Product Pack, you can search for it using the License List.
4. In the search results page, click **Readme** to download and review the Readme file for download instructions and product information.
5. After you review the Readme, choose the appropriate Media Pack from the search results to download the individual zip files. Follow the Download Notes instructions in this page. Once you download and extract the contents of the required zip files, proceed with the installation of the software.

Note: Print the page with the list of downloadable files. It contains a list of part numbers and their corresponding descriptions that you may need to reference during the installation process.

See Also: *Frequently Asked Questions* section on the Oracle E-Delivery Web site for more information about Media Packs

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. If the downloaded file has the `zip` extension, use the following command to extract the content:

```
unzip file_name.zip
```

If the downloaded file has the `cpio.gz` extension, use the following command:

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

This command creates files with names similar to the following:

```
filename.cpio
```

To extract the installation files, enter a command similar to the following:

```
$ cpio -idcmv < filename.cpio
```

Note: Refer to the download page for information about the correct options to use with the `cpio` command.

Some browsers uncompress files while downloading them, but leave the `.gz` file extension.

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

When you have extracted all of the required installation files, refer to the "[Installing the Oracle Database Software](#)" section.

Copying the Software to the Hard Disk

Before installing Oracle Database, you might want to copy the software to the hard disk. This enables the installation process to run a bit faster. Before copying the installation media content to the hard disk, you must mount the disk. The following sections describe to mount disk and copy its content to the hard disk.

Mounting Disks

On most Linux systems, the disk mounts automatically when you insert it into the installation media. If the disk does not mount automatically, then follow these steps to mount it:

1. If necessary, enter a command similar to one of the following to eject the currently mounted disk, then remove it from the drive:

- Asianux, Oracle Enterprise Linux, and Red Hat Enterprise Linux:

```
$ sudo eject /mnt/dvd
```

- SUSE Linux Enterprise Server:

```
# eject /media/dvd
```

In these examples, `/mnt/dvd` and `/media/dvd` are the mount point directories for the installation media.

2. Insert the appropriate installation media into the disk drive.
3. To verify if the disk is mounted automatically, enter one of the following commands depending on the platform:
 - Asianux, Oracle Enterprise Linux, and Red Hat Enterprise Linux:

```
# ls /mnt/dvd
```
 - SUSE Linux Enterprise Server:

```
# ls /media/dvd
```
4. Before running the following command, ensure that the `/mnt/dvd` directory exists on Red Hat Enterprise Linux. If not, create the `/mnt/dvd` as required, to mount the installation media.

If this command fails to display the contents of the installation media, enter a command similar to the following to mount it, depending on the platform:

- Asianux, Oracle Enterprise Linux, and Red Hat Enterprise Linux:

```
# mount -t iso9660 /dev/dvd /mnt/dvd
```
- SUSE Linux Enterprise Server:

```
# mount -t iso9660 /dev/dvd /media/dvd
```

In these examples, `/mnt/dvd` and `/media/dvd` are the mount point directories for the installation media.

5. If Oracle Universal Installer is displaying the Disk Location dialog box, enter the disk mount point directory path, for example:

```
/mnt/dvd
```

To continue, go to one of the following sections:

- If you want to copy software to a hard disk, refer to the ["Copying the Oracle Database Software to a Hard Disk"](#) section.
- If you want to install the software from the installation media, refer to the ["Installing the Oracle Database Software"](#) section.

Copying the Oracle Database Software to a Hard Disk

Note: If the system does not have a installation media, you can copy the software from the disk 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 installation media to a hard disk:

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

```
$ mkdir OraDb11g
```


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

```
$ cd OraDb11g
```

3. Mount the disk, if it is not already mounted.

Some platforms automatically mount the disk when you insert it into the drive. If the disk does not mount automatically, refer to "[Mounting Disks](#)" section on page 4-5 for platform-specific information about mounting it.

4. Copy the contents of the mounted disk to the corresponding new subdirectory as follows:

```
$ cp -R /directory_path OraDb11g
```

In this example, */directory_path* is the disk mount point directory.

5. If necessary, mount the next disk and repeat step 4.

Database Security Options

During installation, you are prompted to select a database security configuration. The Secure Configuration option configures the database with database auditing options, and password policy and expiration settings.

For new database installations, the default configuration for Oracle Database 11g Release 2 (11.2) includes the Secure Configuration option. If you want to disable these enhanced security controls, then you can uncheck the **Assert all new security settings** check box in the Specify Configuration Option screen that appears during installation. Oracle Database is then installed with default options for Oracle Database 10g Release 2. You can enable or disable auditing or password security settings, or revert to a previous security setting. After installation, you can enable or disable the security configuration by starting Database Configuration Assistant through the command line interface. Perform the following:

To Enable Security Configuration:

```
dbca -silent -configureDatabase -sourceDB SID -disableSecurityConfiguration NONE
-enableSecurityConfiguration true
```

To Disable Security Configuration:

```
dbca -silent -configureDatabase -sourceDB SID -disableSecurityConfiguration
[ALL|PASSWORD_PROFILE] -enableSecurityConfiguration false
```

Where *SID* is the system identifier.

For database upgrades, the upgraded database retains your existing database security configuration, to ensure compatibility with existing applications.

Starting with Oracle Database 11g Release 2 (11.2), the Security Enhanced Linux (SE Linux) feature is supported for Oracle Enterprise Linux 4.0, Red Hat Enterprise Linux 4.0, Oracle Enterprise Linux 5.0, and Red Hat Enterprise Linux 5.0.

Note:

- Oracle strongly recommends configuring your database with the Secure Configuration option either during installation, or after installation using Database Configuration Assistant.
 - Database Vault is an enhanced Security feature. If it is installed with the database, then you cannot change the Secure Configuration using Database Configuration Assistant option.
-
-

Installing the Oracle Database Software

In most cases, you use the graphical user interface (GUI) provided by Oracle Universal Installer to install Oracle Database. The instructions in this section explain how to run the Oracle Universal Installer GUI to perform most database installations.

See Also:

- If you plan to use Oracle Restart or Automatic Storage Management, you must install the grid infrastructure before you install the database. For information on installing grid infrastructure refer to ["Installing and Configuring Grid Infrastructure for a Standalone Server"](#) on page 3-14
- [Appendix A](#) if you want to install Oracle Database by using the silent or response file installation method, without the GUI. This method is useful if you must perform multiple installations of Oracle Database. This appendix covers other advanced installation topics as well.

Running Oracle Universal Installer

This section describes the Typical Installation as a default setting. For any type of installation process, start Oracle Universal Installer and install the software, as follows:

1. Log on as the Oracle software owner user (typically, `oracle`) to the computer on which you want to install Oracle components.
2. If you are installing the software from installation media, mount the disk if it is not already mounted.

If the disk does not mount automatically, refer to ["Mounting Disks"](#) section on page 4-5 for platform-specific information about mounting it.

Some platforms automatically mount the disk when you insert the installation media into the drive.

3. To start Oracle Universal Installer, complete one of the following steps depending on the location of the installation files:

Note: Start Oracle Universal Installer from the terminal session where you logged in as the `oracle` user and set the user's environment.

See Also: ["Configuring the oracle User's Environment"](#) for information about setting the `oracle` user's environment

- If the installation files are on installation media, enter commands similar to the following, where *directory_path* is the path of the database directory on the installation media:

```
$ /directory_path/runInstaller
```

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

```
$ ./runInstaller
```

If Oracle Universal Installer is not displayed, refer to "[X Window Display Errors](#)" and "[Remote Terminal Installation Error](#)" for information about troubleshooting.

4. Use the following guidelines to complete the installation:

- Do not install Oracle Database 11g Release 2 (11.2) software into an existing Oracle home.
- Follow the instructions displayed on the Oracle Universal Installer screens. If you need additional information, click **Help**.

See Also: "[Reviewing Accounts and Passwords](#)" on page 6-5 for details on password guidelines

- Do not modify the Java Runtime Environment (JRE) except by using a patch provided by Oracle Support Services. Oracle Universal Installer automatically installs the Oracle-supplied version of the JRE. This version is required to run Oracle Universal Installer and several Oracle assistants.
- If errors are displayed while installing the software, refer to [Appendix G](#) for information about troubleshooting.
- If you chose an installation type that runs Oracle Database Configuration Assistant and Oracle Net Configuration Assistant in interactive mode, then you must provide detailed information about configuring the database and network.

If you need assistance when using the Oracle Database Configuration Assistant or Oracle Net Configuration Assistant in interactive mode, click **Help** on any screen.

Note: If you chose a default installation, Oracle Database Configuration Assistant and Oracle Net Configuration Assistant do not run interactively.

5. When the configuration assistant tasks are complete click finish, click **Exit**, then click **Yes** to exit from Oracle Universal Installer.
6. When Oracle Universal 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
```

7. See [Chapter 5](#) for information about tasks that you must complete after you have installed Oracle Database.

The following table lists the various screens displayed during an Enterprise Edition installation for Oracle Database 11g Release 2 (11.2):

Screen	Action
Configure Security Updates	<p>Enter your e-mail address, preferably your My Oracle Support (formerly Oracle<i>MetaLink</i>) e-mail address or user name in the Email field.</p> <p>You can select the I wish to receive security updates via My Oracle Support check box to receive security updates.</p> <p>Enter your My Oracle Support (formerly Oracle<i>MetaLink</i>) password in the My Oracle Support Password field.</p> <p>Click Next.</p>
Select Installation Option	<p>Select one of the following install options, then click Next:</p> <ul style="list-style-type: none"> ■ Create and Configure a Database: This option creates a new database along with sample schemas. ■ Install Database Software Only: This option only installs the database binaries. To configure database, you must run Oracle Database Configuration Assistant after the software installation. ■ Upgrade an Existing Database: This option installs the software binaries in a new Oracle home. At the end of the installation, you can upgrade the existing database.
System Class	<p>Select the type of system for installing the database, click Next.</p> <ul style="list-style-type: none"> ■ Desktop Class: Choose this option if you are installing on a laptop or desktop class system. This option includes a starter database and allows minimal configuration. This option is designed for those who want to get up and running with the database quickly. See Also: "Setting the ORACLE_HOSTNAME Environment Variable" ■ Server Class: Choose this option if you are installing on a server class system, such as what you would use when deploying Oracle in a production data center. This option allows for more advanced configuration options. Advanced configuration options available using this option include Oracle RAC, Automatic Storage Management, backup and recovery configuration, integration with Enterprise Manager Grid Control, and more fine-grained memory tuning, among many others.
Grid Installation Options	<p>Select the type of database installation you want to perform, click Next.</p> <ul style="list-style-type: none"> ■ Single instance database installation: This option installs the database and the listener. ■ Real Application Clusters database installation: This option installs Oracle Real Application Clusters and the listener.

Screen	Action
Select Install Type	<p>Select one of the following, then click Next:</p> <ul style="list-style-type: none"> ■ Typical Installation: This installation method is selected by default. It lets you quickly install Oracle Database using minimal input. It installs the software and optionally creates a general-purpose database using the information that you specify on this screen. ■ Advanced Installation: This installation method enables to perform more complex installations, such as creating individual passwords for different accounts, creating specific types of starter databases (for example, for transaction processing or data warehouse systems), using different language groups, specifying e-mail notifications, and so on.
Select Product Language	<p>This option enables you to select the language in which you want to run the product.</p> <p>Select the product Language from the Available Languages list, transfer it to the Selected Languages list. Click Next.</p>
Select Database Edition	<p>Select Enterprise Edition, Standard Edition, or Standard Edition One. Click Next.</p> <p>If you click Select Options, then based on your selection you can enable or disable components from the components list. This screen enables you to customize the database. The components available in this screen are:</p> <ul style="list-style-type: none"> ■ Oracle Partitioning ■ Oracle OLAP ■ Oracle Label Security ■ Oracle Data Mining RDBMS Files ■ Oracle Database Vault option ■ Oracle Real Application Testing <p>Click OK to continue.</p> <p>Note: The Select Options button is enabled only if you select the Enterprise Edition installation type.</p>
Specify Installation Location	<p>The Oracle base path appears by default. You can change the path based on your requirement.</p> <p>In the Software Location section, accept the default values or enter the Oracle home name and directory path in which you want to install Oracle components. The directory path should not contain spaces. Click Next.</p> <p>Note: This screen is available only with Advanced Installation.</p> <p>See Also: "Naming Directories" for information on directory naming conventions</p>
Create Inventory	<p>You are prompted by the installer to specify the directory path for central inventory the first time you install any Oracle software on your computer. Click Next.</p> <p>Note: By default, the Oracle Inventory directory is not installed under the Oracle Base directory. This is because all Oracle software installations share a common Oracle Inventory, so there is only one Oracle Inventory for all users, whereas there is a separate Oracle Base for each user.</p>

Screen	Action
Select Configuration Type	<p>Select one of the following, click Next:</p> <ul style="list-style-type: none"><li data-bbox="716 275 1308 352">■ General Purpose / Transaction Processing: This is a starter database designed for general usage or transaction-heavy applications.<li data-bbox="716 365 1349 422">■ Data Warehousing: A starter database optimized to run Data Warehousing applications. <p>See the online Help provided by either Oracle Universal Installer or Oracle Database Configuration Assistant for a description of these preconfigured database types.</p>
Specify Database Identifiers	<p>Specify the following information, then click Next:</p> <p>Database Naming</p> <p>Specify the Global Database Name using the following syntax:</p> <p><i>database_name.domain</i></p> <p>where:</p> <ul style="list-style-type: none"><li data-bbox="716 783 1308 861">■ <i>database_name</i> is the name of the database. It can contain no more than 8 characters (alphanumeric, underscore (_), dollar (\$), and pound (#)).<li data-bbox="716 873 1349 951">■ <i>domain</i> is the domain used for the database. It can contain no more than 128 characters (alphanumeric, underscore (_), and pound (#)), inclusive of all periods. <p>For example:</p> <p><i>sales.us.example.com</i></p> <p>When you enter the Global Database Name, Oracle Universal Installer automatically populates the SID field with the database name, but you can change this SID to another name. The SID can have no more than 64 alphanumeric characters.</p> <p>See Also: "Setting the ORACLE_HOSTNAME Environment Variable"</p>

Screen	Action
Specify Configuration Options	<p>Specify the following configuration details, click Next:</p> <p>Memory</p> <p>Enable Automatic Memory Management option is selected by default. This option enables the database to automatically distribute memory between SGA and PGA. If you deselect this option, then the SGA and PGA must be sized manually.</p> <p>Character Sets</p> <p>This option enables you to store the character data in the database in one of the following methods:</p> <ul style="list-style-type: none"> Use the default: This option makes use of the operating system language settings. Use Unicode: This option enables you to store multiple language groups Choose from the list of character sets: This option enables the Select Database Character Set drop down list. <p>See Also:</p> <ul style="list-style-type: none"> "Selecting the Database Character Set" <i>Oracle Database Globalization Support Guide</i> for information on choosing a character set <p>Security</p> <p>The Assert all new security settings option is selected by default. This setting includes enabling auditing and using new password profile.</p> <p>Note: Oracle recommends that you use the default settings.</p> <p>Sample Schema</p> <p>The Create database with sample schema option is not selected by default. However, you can select the option, if you want to create the starter database with sample schema.</p>
Specify Management Option	<p>Select one of the following, then click Next:</p> <ul style="list-style-type: none"> Use an Existing Grid Control for Database Management: This option is useful if you have Oracle Enterprise Manager installed. Use Database Control for Database Management: This option enables you to manage Oracle Database locally. Optionally, select Enable Email Notifications and enter the outgoing SMTP server and e-mail address.
Specify Database Storage Options	<p>Select one of the following, click Next.</p> <ul style="list-style-type: none"> File System: Specify the database file location. Automatic Storage Management: Specify a password for the ASMSNMP user. <p>Note: Installing Oracle data files on an Oracle ACFS file system is not supported. Oracle recommends that these data files are installed in Oracle ASM disk groups.</p>
Specify Backup and Recovery Options	<p>Select one of the following, click Next.</p> <ul style="list-style-type: none"> Disable automated backups Enable automated backups: If you select this option, then the backup job will use a specified recovery storage area.

Screen	Action
Specify Database Schema Passwords	<p>Enter and confirm passwords for the privileged database accounts, click Next.</p> <p>Note: Optionally, you can use the same password for all accounts. However, Oracle recommends that you specify a different password for each account. You must remember the passwords that you specify.</p> <p>Refer to "Unlocking and Changing Passwords" on page 6-9 for information on password guidelines.</p>
Privileged Operating System Groups	<p>The operating system groups are selected by default. You can also manually select the OSDBA and OSOPER groups.</p> <p>Click Next.</p> <p>See Also: "Creating Required Operating System Groups and Users" for information on operating system groups and users</p>
Perform Prerequisite Checks	<p>This option checks if the minimum system requirements to carry out the database installation are met.</p> <p>If you click Check Again, then you can run the prerequisite check again to see if the minimum requirements are met to carry on with the database installation.</p> <p>Click Fix & Check Again, if you want the installer to fix the problem and check the system requirements once more.</p> <p>Note: The Fix & Check Again option generates a script that you need to run as the <code>root</code> user. This generated script sets some of the system parameters to Oracle-recommended values. Oracle recommends that you do not modify the contents of this script. Refer to "Installation Fixup Scripts" for more information on fixup scripts.</p> <p>To get a list of failed requirements, select ShowFailed from the list. To get a list of all the prerequisite checks run by the OUI, select Show All. To get a list of the prerequisite checks that are successful, select Show Succeeded.</p> <p>Note: You can check Ignore All, to ignore all the errors and proceed with the database installation.</p> <p>See Also: Chapter 2, "Oracle Database Preinstallation Requirements" for information on the system requirements</p>
Summary	<p>Review the information displayed on this screen, Click Finish</p> <p>Note: Starting with Oracle Database 11g Release 2 (11.2), you can save all the installation steps into a response file by clicking Save Response File. Later, this file can be used for a silent installation.</p>
Install product	<p>This screen states the progress of a database installation. After the database is installed, you are prompted to execute some root configuration script for new inventory as a root user. Click Next.</p> <p>This screen then displays the status information for the configuration assistants that configure the software and create a database.</p> <p>Finally, a message is displayed at the end of Database Configuration Assistant process, and click OK.</p> <p>Execute the <code>root.sh</code> script as the <code>root</code> user to complete the installation, and click OK.</p>
Finish	<p>This screen is shown automatically when all the configuration tools are successful.</p> <p>Click Close.</p>

Caution: After installation is complete, do not manually remove, or run cron jobs that remove `/tmp/.oracle` or `/var/tmp/.oracle` directories or their files while Oracle software is running. If you remove these files, then Oracle software can encounter intermittent hangs. Oracle Restart installations will fail with the following error:

CRS-0184: Cannot communicate with the CRS daemon.

Installing Oracle Database Examples

If you plan to use the following products or features, then download and install the products from the Oracle Database Examples media:

- Oracle JDBC Development Drivers
- Oracle Database Examples
- Oracle Text Knowledge Base
- Various Oracle product demonstrations

For information about installing software and various Oracle product demonstrations from the Oracle Database Examples media, refer to *Oracle Database Examples Installation Guide*.

Oracle Database Postinstallation Tasks

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

- [Required Postinstallation Tasks](#)
- [Recommended Postinstallation Tasks](#)
- [Required Product-Specific Postinstallation Tasks](#)
- [Postinstallation Tasks for SQL Developer](#)

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

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

Note: This chapter describes basic configuration only. Refer to *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems*, *Oracle Database Administrator's Guide* and product-specific administration and tuning guides for more detailed configuration and tuning information.

See Also: "Post-installation Database Configuration" section in *Oracle Configuration Manager Installation and Administration Guide* for information on postinstallation tasks for Oracle Configuration Manager

Required Postinstallation Tasks

Perform the task described in the following section after completing an installation.

Downloading and Installing Patches

Check the My Oracle Support (formerly *OracleMetaLink*) Web site for required patches for the installation.

To download required patches:

1. Use a Web browser to view the My Oracle Support (formerly *OracleMetaLink*) Web site:

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

2. Log in to My Oracle Support (formerly *OracleMetaLink*).

Note: If you are not an My Oracle Support (formerly *OracleMetaLink*) registered user, click **Register here** and follow the registration instructions.

3. On the main My Oracle Support (formerly *OracleMetaLink*) page, click **Patches and Updates**.
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.
6. Find the latest patch set for Oracle Database using My Oracle Support (formerly *OracleMetaLink*).
7. From the list of available patches, select a patch to download.

Patch sets for Oracle databases are identified as *x.x.x* PATCH SET FOR ORACLE DATABASE SERVER.
8. Review the README file before proceeding with the download.

Each patch has a README file with installation requirements and instructions. Some patches install with Oracle Universal Installer; others require special procedures. Oracle recommends that you always read the README file before proceeding.
9. Download and install the patch.

Recommended Postinstallation Tasks

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

- [Creating a Backup of the root.sh Script](#)
- [Configuring New or Upgraded Databases](#)
- [Creating and Configuring Additional Operating System Accounts](#)
- [Setting the NLS_LANG Environment Variable](#)
- [Generating the Client Static Library](#)
- [Guidelines for Setting Semaphore Parameters](#)
- [Create a Fast Recovery Area Disk Group](#)

Creating a Backup of the root.sh Script

Oracle recommends that you back up the `root.sh` script after you complete an installation. If you install other products in the same Oracle home directory, then Oracle Universal Installer updates the contents of the existing `root.sh` script during

the installation. If you require information contained in the original `root.sh` script, then you can recover it from the backed up `root.sh` file.

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.

See Also: *Oracle Database Upgrade Guide* for more information about database upgrade

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> @?/rdbms/admin/utlrp.sql
```

Creating and Configuring Additional Operating System Accounts

If required, create additional operating system accounts. Users must be members of the `OSDBA` or `OSOPER` groups to connect to the database with administrator privileges.

Configuring the Accounts of Oracle Users

Update the startup files of the `oracle` user and the operating system accounts of Oracle users, specifying the appropriate environment variables in the environment file.

For the Bourne, Bash, or Korn shell, add the environment variables to the `.profile` file, or the `.bash_profile` file for the Bash shell on Red Hat Enterprise Linux.

For the C shell, add the environment variables to the `.login` file.

Note: You can use the `oraenv` or `coraenv` script to ensure that Oracle user accounts are updated.

Setting the NLS_LANG Environment Variable

NLS_LANG is an environment variable that specifies the locale behavior for Oracle software. This variable sets the language and territory used by the client application and the database server. It also declares the character set of the client, which is the character set of data entered or displayed by an Oracle client program, such as SQL*Plus.

See Also: [Appendix F, "Configuring Oracle Database Globalization Support"](#) for more information about the NLS_LANG environment variable

Generating the Client Static Library

The client static library (`libclntst11.a`) is not generated during installation. If you want to link the 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/11.2.0/dbhome_1
$ export ORACLE_HOME
```

- C shell:

```
% setenv ORACLE_HOME /u01/app/oracle/product/11.2.0/dbhome_1
```

3. Enter the following command:

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

Guidelines for Setting Semaphore Parameters

Refer to the following guidelines only if the default semaphore parameter values are too low to accommodate all Oracle processes:

Note: Oracle recommends that you refer to the operating system documentation for more information about setting semaphore parameters.

1. Calculate the minimum total semaphore requirements using the following formula:

```
sum (process parameters of all database instances on the system) + system and
other application requirements
```

2. Set `semnms` (total semaphores systemwide) to this total.
3. Set `semmsl` (semaphores per set) to 256.
4. Set `semmni` (total semaphores sets) to `semnms / semmsl` rounded up to the nearest multiple of 1024.

Create a Fast Recovery Area Disk Group

During installation, by default you can create one disk group. If you plan to add an Oracle Database for a standalone server, then you should create the fast recovery area for database files.

About the Fast Recovery Area and the Fast Recovery Area Disk Group

The fast recovery area is a unified storage location for all Oracle Database files related to recovery. Database administrators can define the `DB_RECOVERY_FILE_DEST` parameter to the path for the fast recovery area to enable on-disk backups, and rapid recovery of data. Enabling rapid backups for recent data can reduce requests to system administrators to retrieve backup tapes for recovery operations.

When you enable fast recovery in the `init.ora` file, all RMAN backups, archive logs, control file automatic backups, and database copies are written to the fast recovery area. RMAN automatically manages files in the fast recovery area by deleting obsolete backups and archive files no longer required for recovery.

Oracle recommends that you create a fast recovery area disk group. Oracle Clusterware files and Oracle Database files can be placed on the same disk group, and you can also place fast recovery files in the same disk group. However, Oracle recommends that you create a separate fast recovery disk group to reduce storage device contention.

The fast recovery area is enabled by setting `DB_RECOVERY_FILE_DEST`. The size of the fast recovery area is set with `DB_RECOVERY_FILE_DEST_SIZE`. As a general rule, the larger the fast recovery area, the more useful it becomes. For ease of use, Oracle recommends that you create a fast recovery area disk group on storage devices that can contain at least three days of recovery information. Ideally, the fast recovery area should be large enough to hold a copy of all of your datafiles and control files, the online redo logs, and the archived redo log files needed to recover your database using the datafile backups kept under your retention policy.

Multiple databases can use the same fast recovery area. For example, assume you have created one fast recovery area disk group on disks with 150 GB of storage, shared by three different databases. You can set the size of the fast recovery for each database depending on the importance of each database. For example, if `database1` is your least important database, `database2` is of greater importance and `database3` is of greatest importance, then you can set different `DB_RECOVERY_FILE_DEST_SIZE` settings for each database to meet your retention target for each database: 30 GB for `database1`, 50GB for `database2`, and 70GB for `database3`.

Creating the Fast Recovery Area Disk Group

To create a fast recovery file disk group:

1. Navigate to the Grid home bin directory, and start ASM Configuration Assistant (ASMCA). For example:

```
$ cd /u01/grid/bin
$ ./asmca
```

2. ASMCA opens at the Disk Groups tab. Click **Create** to create a new disk group.
3. The Create Disk Groups window opens.

In the Disk Group Name field, enter a descriptive name for the fast recovery area group. For example: `FRA`.

In the Redundancy section, select the level of redundancy you want to use.

In the Select Member Disks field, select eligible disks to be added to the fast recovery area, and click **OK**.

4. The Diskgroup Creation window opens to inform you when disk group creation is complete. Click **OK**.
5. Click **Exit**.

See Also: ■

- "Setting the Fast Recovery Area Location and Initial Size" section in *Oracle Database Backup and Recovery User's Guide*
- *Oracle Database Storage Administrator's Guide*

Required Product-Specific Postinstallation Tasks

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

- [Configuring Oracle Net Services](#)
- [Configuring Oracle Label Security](#)
- [Configuring Oracle Database Vault](#)
- [Configuring Oracle Messaging Gateway](#)
- [Configuring Oracle Precompilers](#)
- [Configuring Secure Sockets Layer](#)
- [Installing Oracle Text Supplied Knowledge Bases](#)
- [Configuring or Reinstalling Oracle XML DB](#)
- [Configuring and Using Direct NFS Client](#)

Note: You need only perform postinstallation tasks for products that you intend to use.

Configuring Oracle Net Services

If you have an earlier 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 earlier release to the corresponding files for the new release.

The following sections describe about how to configure the Oracle Net Services:

- [Modifying the listener.ora File](#)
- [Modifying the tnsnames.ora File](#)

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.

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.

If you have referenced the previous Oracle home directory names in the static listener information, then these directory names must be modified before the `listener.ora` file can be used in the 11.2 environment.

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 earlier 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: *Oracle Label Security Administrator's Guide* for more information about Oracle Label Security enabled with Oracle Internet Directory

Configuring Oracle Database Vault

If you have installed Oracle Database Vault, then you must register it in a database. Ensure that you create the Database Vault Owner and, optionally, Database Vault Account Manager administrative accounts before you can use it.

See Also: *Oracle Database Vault Administrator's Guide* for more information on registering Oracle Database Vault

Configuring Oracle Messaging Gateway

To configure Oracle Messaging Gateway, refer to the section about Messaging Gateway in *Oracle Streams Advanced Queuing User's Guide*. When following the

instructions listed in that manual, refer to this section for additional 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 = "LD_LIBRARY_PATH=/oracle_home/jdk/jre/lib/i386:/oracle_
home/jdk/jre/lib/i386/server:/oracle_home/lib")
  (ORACLE_HOME = oracle_home)
  (PROGRAM = extproc)
)
```

In this example:

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

In the settings for 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.

- `oracle_home` is the path of the Oracle home directory.
- `extproc` is the external procedure agent executable file

The following example shows a sample `listener.ora` file:

```
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = PLSExtProc)
      (ORACLE_HOME = /u01/app/oracle/product/11.2.0/dbhome_1)
      (PROGRAM = extproc)
    )
    (SID_DESC =
      (SID_NAME = mgwextproc)
      (ENVS = "LD_LIBRARY_PATH =/u01/app/oracle/product/11.2.0/dbhome_
1/jdk/jre/lib/i386:/u01/app/oracle/product/11.2.0/dbhome_
1/jdk/jre/lib/i386/server:/u01/app/oracle/product/11.2.0/dbhome_1/lib")
      (ORACLE_HOME = /u01/app/oracle/product/11.2.0/dbhome_1)
      (PROGRAM = extproc)
    )
  )
)
```

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)))
```

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
JRE runtime	<code>\$ORACLE_HOME/jdk/jre/lib/rt.jar</code>
Oracle JDBC	<code>\$ORACLE_HOME/jdbc/lib/ojdbc5.jar</code>
Oracle internationalization	<code>\$ORACLE_HOME/jlib/orai18n.jar</code>
SQLJ	<code>\$ORACLE_HOME/sqlj/lib/runtime12.jar</code>
JMS Interface	<code>\$ORACLE_HOME/rdbms/jlib/jmscommon.jar</code>
Oracle JMS implementation	<code>\$ORACLE_HOME/rdbms/jlib/aqapi.jar</code>
Java Transaction API	<code>\$ORACLE_HOME/jlib/jta.jar</code>

Note: All the lines in the `mgw.ora` file should consist of less than 1024 characters.

Configuring Oracle Precompilers

This section describes postinstallation tasks for Oracle precompilers:

- [Configuring Pro*C/C++](#)
- [Configuring Pro*FORTRAN](#)

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 5–1 shows the default directories and the appropriate command to verify the path setting of the compiler.

Table 5–1 C/C++ Compiler Directory

Path	Command
/usr/bin	\$ which gcc
/opt/intel_cce_80/bin/icc	\$ which icc

Configuring Pro*FORTRAN

Verify that the PATH environment variable setting includes the directory that contains the FORTRAN compiler executable. You can verify the path setting by using the `which xlf` command. The path for the FORTRAN executable is `/usr/bin`.

Configuring Secure Sockets Layer

Oracle highly recommends you configure and use a Secure Sockets Layer (SSL) to ensure that passwords and other sensitive data are not transmitted in clear text in HTTP requests.

See Also:

- "Using SSL" and "Enabling SSL" in *Oracle Database Advanced Security Administrator's Guide* for more information on configuring and using SSL
- "SSL Usage Issues" in *Oracle Database Advanced Security Administrator's Guide* for more information on SSL usage issues

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, then you can install two supplied knowledge bases (English and French).

See Also:

- *Oracle Database Examples Installation Guide*
- *Oracle Text Reference* for information about creating and extending knowledge bases, such as extending the supplied knowledge bases to accommodate your requirements, or creating your own knowledge bases in languages other than English and French

Configuring or Reinstalling Oracle XML DB

Refer to *Oracle XML DB Developer's Guide* for more information about the following tasks:

- Reinstalling Oracle XML DB
- Configuring or customizing the Oracle XML DB tablespace
- Configuring FTP, HTTP/WebDAV port numbers

See Also: Appendix A of *Oracle XML DB Developer's Guide*

Configuring and Using Direct NFS Client

Many Oracle products and options must be configured before you use them for the first time. Before using individual Oracle products or options, refer to the appropriate manual in the product documentation library.

See Also: ■

- Chapter 4, "Configuration Tasks When Installing from the Database" in *Oracle Application Express Installation Guide*
- ["Using NAS Devices"](#)

Direct NFS Client

With Oracle Database 11g, instead of using the operating system kernel NFS client, you can configure Oracle Database to access NFS V3 servers directly using an Oracle internal Direct NFS client. If Oracle Database is unable to open an NFS server using Direct NFS, then Oracle Database uses the platform operating system kernel NFS client. In this case, the kernel NFS mount options must be set up as defined in ["Checking NFS Buffer Size Parameters"](#). Additionally, an informational message will be logged into the Oracle alert and trace files indicating that Direct NFS could not be established.

The Oracle files resident on the NFS server that are served by the Direct NFS Client are also accessible through the operating system kernel NFS client. The usual considerations for maintaining integrity of the Oracle files apply in this situation.

Some NFS file servers require NFS clients to connect using reserved ports. If your filer is running with reserved port checking, then you must disable it for Direct NFS to operate. To disable reserved port checking, consult your NFS file server documentation.

Direct NFS can use up to four network paths defined for an NFS server. The Direct NFS client performs load balancing across all specified paths. If a specified path fails, then Direct NFS reissues I/O commands over any remaining paths.

The following sections elaborate on enabling, disabling, checking the buffer size for a Direct NFS Client:

- [Enabling Direct NFS Client](#)
- [Disabling Direct NFS Client](#)
- [Checking NFS Buffer Size Parameters](#)

Enabling Direct NFS Client By default Direct NFS will attempt to serve mount entries found in `/etc/mtab`. No other configuration is required. You can use `oranfstab` to specify additional Oracle specific options to Direct NFS. For example, you can use `oranfstab` to specify additional paths for a mount point.

A new Oracle specific file `oranfstab` can be added to either `/etc` or to `$ORACLE_HOME/dbs`. When `oranfstab` is placed in `$ORACLE_HOME/dbs`, its entries are specific to a single database. However, when `oranfstab` is placed in `/etc`, then it is global to all Oracle databases, and hence can contain mount points for all Oracle databases.

Note: Direct NFS will not work and fall back to traditional kernel NFS path if the backend NFS server does not support a write size (`wtmax`) of 32768 or larger.

Direct NFS determines mount point settings to NFS storage devices based on the configurations in `/etc/mtab`. Direct NFS looks for the mount point entries in the following order:

1. `$ORACLE_HOME/dbs/oranfstab`
2. `/etc/oranfstab`
3. `/etc/mtab`

It uses the first matched entry as the mount point.

In all cases, Oracle requires that mount points be mounted by the kernel NFS system even when being served through Direct NFS.

Complete the following procedure to enable Direct NFS:

1. You can optionally create an `oranfstab` file with the following attributes for each NFS server to be accessed using Direct NFS:
 - **Server:** The NFS server name.
 - **Path:** Up to four network paths to the NFS server, specified either by IP address, or by name, as displayed using the `ifconfig` command on the filer.
 - **Local:** Up to four local paths on the database host, specified by IP address or by name, as displayed using the `ifconfig` command run on the database host.
 - **Export:** The exported path from the NFS server.
 - **Mount:** The corresponding local mount point for the exported volume.
 - **Dontroute:** Specifies that outgoing messages should not be routed by the operating system, but sent using the IP address they are bound to.
 - **mnt_timeout:** Specifies (in seconds) the time for which Direct NFS client should wait for a successful mount before timing out. This parameter is optional and the default timeout is 10 minutes.

Note: The parameters `local` and `dontroute` are available from patchset 11.1.0.7 onwards.

The following is an example of an `oranfstab` file with two NFS server entries:

```
server: MyDataServer1
local: 145.34.45.12
path: 132.34.35.12
local: 132.34.45.13
path: 132.34.35.13
dontroute
export: /vol/oradata1 mount: /mnt/oradata1

server: MyDataServer2
local: LocalPath1
path: NfsPath1
local: LocalPath2
path: NfsPath2
local: LocalPath3
path: NfsPath3
local: LocalPath4
path: NfsPath4
dontroute
```

```
export: /vol/oradata2 mount: /mnt/oradata2
export: /vol/oradata3 mount: /mnt/oradata3
export: /vol/oradata4 mount: /mnt/oradata4
export: /vol/oradata5 mount: /mnt/oradata5
```

2. Oracle Database uses an ODM library, `libnfsodm11.so`, to enable Direct NFS. To replace the standard ODM library, `$ORACLE_HOME/lib/libodm11.so`, with the ODM NFS library, `libnfsodm11.so`, complete the following steps:

- Change directory to `$ORACLE_HOME/lib`.
- Enter the following commands:

```
cp libodm11.so libodm11.so_stub
ln -s libnfsodm11.so libodm11.so
```

Disabling Direct NFS Client Use one of the following methods to disable the Direct NFS client:

- Remove the `oranfstab` file.
- Restore the stub `libodm11.so` file by reversing the process you completed in [Enabling Direct NFS Client](#).
- Remove the specific NFS server or export paths in the `oranfstab` file.

Note: If you remove an NFS path that Oracle Database is using, then you must restart the database for the change to be effective.

Checking NFS Buffer Size Parameters If you are using NFS, then you must set the values for the NFS buffer size parameters `rsize` and `wsiz` to at least 16384. Oracle recommends that you use the value 32768.

Direct NFS will issue writes at `wtmax` granularity to the NFS server.

For example, if you decide to use `rsize` and `wsiz` buffer settings with the value 32768, then update the `/etc/vfstab` file on each node with an entry similar to the following:

```
nfs_server:/vol/DATA/oradata /home/oracle/netapp nfs\
rw,bg,hard,nointr,rsize=32768,wsiz=32768,tcp,actimeo=0,vers=3,timeo=600
```

Note: Refer to your storage vendor documentation for additional information about mount options.

Postinstallation Tasks for SQL Developer

This section describes tasks that you must complete after you install the software:

- [Migrating User Settings from Release 1.0](#)
- [Migrating Information from Previous Releases](#)
- [Location of User-Related Information](#)

Migrating User Settings from Release 1.0

The first time you start SQL Developer after installing it or after adding any extensions, you are asked if you want to migrate your user settings from a previous

release. (This occurs regardless of whether there was a previous release on your system.)

Note: Migration of user settings is supported only from SQL Developer release 1.0 to release 1.1. It is *not* supported for migration from a pre-release version of 1.1 to release 1.1.

These settings refer to database connections, reports, and certain SQL Developer user preferences that you set in a previous version by clicking **Tools** and then **Preferences**. However, some user preferences are not saved, and you must re-specify these using the new release.

To migrate user settings from SQL Developer release 1.0:

1. Unzip the release 1.1 kit into an empty directory (folder). Do not delete or overwrite the directory into which you unzipped the release 1.0 kit.
2. When you start SQL Developer release 1.1, click **Yes** when asked if you want to migrate settings from a previous release.
3. In the dialog box that is displayed, do *not* accept the default location for the settings. Instead, specify the location of your release 1.0 settings, which might be a folder whose path ends with `sqldeveloper/jdev/system`.

See Also: ["Migrating Information from Previous Releases"](#) on page 5-14 for more information

Migrating Information from Previous Releases

If you have used a previous release of SQL Developer or a pre-release version of the current release, you may want to preserve database connections that you have been using. To preserve database connections, save your existing database connections in an XML file. To save the connections, right-click the Connections node in the Connections Navigator and select **Export Connections**. After you complete the installation described in this guide, you can use those connections by right-clicking the Connections node in the Connections Navigator and selecting **Import Connections**.

If you want to use any user-defined reports or the SQL history from a previous version, see ["Location of User-Related Information"](#) on page 5-14 for information about where these are located. If you want to use any user-defined reports or the SQL history from release 1.0 with both releases 1.0 and 1.1, you must save them before using release 1.1, because release 1.1 modifies the files to a format that is incompatible with release 1.0.

SQL Developer preferences (specified by clicking Tools and then Preferences) from a pre-release version of the current release cannot currently be saved and reused; you must re-specify any desired preferences.

Location of User-Related Information

SQL Developer stores user-related information in several places, with the specific location depending on the operating system and certain environment specifications. User-related information includes user-defined reports, user-defined snippets, SQL Worksheet history, and SQL Developer user preferences.

The user-related information is stored outside the SQL Developer installation directory hierarchy, so that it is preserved if you delete that directory and install a new version.

This information is stored in or under the `SQLDEVELOPER_USER_DIR` location, if defined; otherwise as indicated in the following table.

The table shows the typical default locations (under a directory or in a file) for specific types of resources on different operating systems. (Note the period in the name of any directory or folder named `.sqldeveloper`.)

Table 5–2 Default Locations for User-Related Information

Resource Type	Linux
User-defined reports	<code>~/.sqldeveloper/UserReports.xml</code>
User-defined snippets	<code>~/.sqldeveloper/UserSnippets.xml</code>
SQL history	<code>~/.sqldeveloper/system/</code>
SQL Worksheet archive files	<code>~/.sqldeveloper/tmp/</code>
SQL Developer user preferences	<code>~/.sqldeveloper/system/</code>

SQL Worksheet archive files contain SQL statements that you have entered. These files begin with `sqldev` and then have a random number (for example, `sqldev14356.sql`). If you close SQL Developer with a SQL Worksheet open that contains statements, then you will be prompted to save these files.

To specify a nondefault `SQLDEVELOPER_USER_DIR` location, do either of the following:

- Set the `SQLDEVELOPER_USER_DIR` environment variable to specify another directory path.
- Edit the `sqldeveloper_install\sqldeveloper\sqldeveloper\bin\sqldeveloper.conf` file and substitute the desired directory path for `SQLDEVELOPER_USER_DIR` in the following line:

```
SetUserHomeVariable SQLDEVELOPER_USER_DIR
```

If you want to prevent other users from accessing your user-specific SQL Developer information, you must ensure that the appropriate permissions are set on the directory where that information is stored or on a directory preceding it in the path hierarchy. For example, you may want to ensure that the `~/.sqldeveloper` directory is not world-readable.

Getting Started with Oracle 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:

- [Checking the Installed Oracle Database Contents and Directory Location](#)
- [Logging In to Oracle Enterprise Manager Database Control](#)
- [Managing Automatic Storage Management](#)
- [Accessing Oracle Database with SQL*Plus](#)
- [Accessing Oracle Database with SQL Developer](#)
- [Reviewing Accounts and Passwords](#)
- [Unlocking and Resetting User Passwords](#)
- [Identifying Databases](#)
- [Locating the Server Parameter File](#)
- [Reviewing Tablespaces and Data Files, Redo Log Files, and Control Files](#)

Checking the Installed Oracle Database Contents and Directory Location

You can use Oracle Universal Installer to check the contents and directory location of an Oracle Database installation. To do this, perform the following steps:

1. To start Oracle Universal Installer, run the following command:

```
$ ORACLE_HOME/oui/bin/runInstaller
```
2. Click **Installed Products** to display the Inventory dialog box on the Welcome screen.
3. Select the Oracle Database product from the list to check the installed contents.
4. Click **Details** to find additional information about an installed product.
5. Click **Close** to close the Inventory dialog box.
6. Click **Cancel** to close Oracle Universal Installer, and then click **Yes** to confirm.

Logging In to Oracle Enterprise Manager Database Control

If you configured Oracle Enterprise Manager Database Control (Database Control) during the installation, you can use it to manage the database. Alternatively, you can

use Oracle Enterprise Manager Grid Control to manage the database. To display the Database Control:

1. First check the Database Control status using the following command:

```
emctl status dbconsole
```

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

```
https://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) = 1158
```

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 1158, use the following URL:

```
https://mgmt42:1158/em
```

Oracle Enterprise Manager displays the Database Control login page.

3. Log in to the database using the user name `SYSTEM` and connect as `SYSDBA`. Enterprise Manager displays the Database Home page.

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

Note: You can also log in to the Database Control using the `SYS`, `SYSTEM` or `SYSMAN` accounts or you can grant login privileges to other database users. If you log in as `SYS`, then you must connect as `SYSDBA`.

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 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: The "[Logging In to Oracle Enterprise Manager Database Control](#)" section 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 an 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. In the **E-mail Address** field, specify one or more e-mail addresses.
7. In the Administrator Privilege list, select either **View on all targets**, **Full on all targets**, or **Super Administrator**.
8. You can also select the **Grant SELECT_CATALOG_ROLE** option.
9. Enter the password for this user, and then click **Review**.
10. On the properties page, 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.

Managing Automatic Storage Management

This section provides information about managing an Automatic Storage Management installation. It covers the following topics:

- [Starting and Stopping Automatic Storage Management](#)
- [Automatic Storage Management Utilities](#)

Starting and Stopping Automatic Storage Management

To start and stop Automatic Storage Management, refer to *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems*.

Automatic Storage Management Utilities

To manage Automatic Storage Management, you can use the following tools:

- **asmcmd**: This command-line tool enables you to manage Automatic Storage Management disk group files and directories.
- **ASMCA**: Oracle Automatic Storage Management Configuration Assistant (ASMCA) is an interactive utility that allows you to create a new Automatic Storage Management instance or upgrade existing Automatic Storage Management instances. It also allows you to create and configure disk groups, Automatic Storage Management volumes and Oracle Automatic Storage Management Cluster File System (Oracle ACFS).
- **Oracle Enterprise Manager Grid Control**: If you have Oracle Enterprise Manager installed, you can use Grid Control to manage Automatic Storage Management functions, such as migrating an existing database to Automatic Storage Management, checking the status of the Automatic Storage Management instance, checking the performance of the Automatic Storage Management disk groups, and creating or dropping Automatic Storage Management disk groups.
- **Oracle Enterprise Manager Database Control**: This utility enables you to perform functions similar to Grid Control.
- **SQL*Plus**: You can run commands that are specific to Automatic Storage Management from either of these tools. To connect to an Automatic Storage Management instance, use the same methods that you use to connect to an Oracle database instance.

See Also:

- ["Logging In to Oracle Enterprise Manager Database Control"](#)
- *Oracle Database Utilities* for more information about `asmcmd`
- *Oracle Database Storage Administrator's Guide* for more information about managing your storage with Automatic Storage Management

Accessing Oracle Database with SQL*Plus

To run the SQL and PL/SQL commands to access the Oracle Database, you can use SQL*Plus. This tool enables you to perform the same database management operations, as well as to query, insert, update, or delete data directly in the database.

Note: Before you start SQL*Plus, ensure that all the environment variables, specially `ORACLE_HOME` and `ORACLE_SID`, are set. Refer to, ["Configuring the oracle User's Environment"](#) for more information on setting environment variables.

Use the following command to start SQL*Plus and log in as the `SYS` user, connecting as `SYSDBA`:

```
$ $ORACLE_HOME/bin/sqlplus
SQL> CONNECT SYS as SYSDBA
Enter password: SYS_password
```

For example, to log on as `SYSTEM` using the password `Systempwd1`, you enter:

```
$ $ORACLE_HOME/bin/sqlplus
SQL> CONNECT SYSTEM
Enter password: Systempwd1
```

If you are logging on as `SYS`, you would need to connect as `SYSDBA`:

```
$ $ORACLE_HOME/bin/sqlplus
SQL> CONNECT SYS as SYSDBA
Enter password: SYS_password
```

See Also:

- ["Connecting to the Database with SQL*Plus"](#) in *Oracle Database Administrator's Guide* for more information on accessing Oracle Database using SQL*Plus
- *SQL*Plus User's Guide and Reference*
- *SQL*Plus Quick Reference*

Accessing Oracle Database with SQL Developer

To run the SQL and PL/SQL commands to access Oracle Database, you can use SQL Developer. All SQL and PL/SQL commands are supported as they are passed directly from the SQL Worksheet to the Oracle Database.

Set Up the JDK Path For SQL Developer

Set the following environmental variables to ensure that the correct `jdk` is picked up:

- `$ORACLE_HOME`
- `$JAVA_HOME=$ORACLE_HOME/jdk`
- `$PATH=$JAVA_HOME/bin/:$PATH`

To start SQL Developer on which the Sun Java SDK release 1.5 is installed, use the following commands:

- Change to `$ORACLE_HOME/sqldeveloper`.
- Run `$./sqldeveloper.sh`.
- Right-Click Connections. In the dialog box, enter a Connection name, username, password, and for the host string, the name of the database to which you want to connect and click Connect.

Once connected, you can view, create, modify, and delete the database objects using the Connection Navigator or issue any SQL or PL/SQL command using a SQL Worksheet (From the **Tools** menu, select **SQL Worksheet**).

SQL*Plus commands have to be interpreted by the SQL Worksheet before being passed to the database. The SQL Worksheet currently supports a number of SQL*Plus commands. SQL*Plus commands which are not supported by the SQL Worksheet are ignored and are not sent to the Oracle Database.

See Also:

"SQL*Plus Statements Supported and Not Supported in SQL Worksheet" in *Oracle Database SQL Developer User's Guide*

Reviewing 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 accounts, you must unlock them and reset their passwords. [Table 6–1](#) describes these accounts and lists their user names and default passwords.

See Also: The "[Unlocking and Resetting User Passwords](#)" section for information about unlocking and resetting passwords

Note: Use the Oracle Enterprise Manager Database Control to view the complete list of database accounts.

Table 6–1 Database Accounts

User Name	Description	See Also
ANONYMOUS	Enables HTTP access to Oracle XML DB.	None
APEX_030200	The account owns the Application Express schema and metadata.	<i>Oracle Application Express Application Builder User's Guide</i>
APEX_PUBLIC_USER	The minimally privileged account used for Application Express configuration with Oracle HTTP Server and <code>mod_plsql</code> .	<i>Oracle Application Express Application Builder User's Guide</i>
APPQOSSYS	Used for storing/managing all data and metadata required by Oracle Quality of Service Management.	None

Table 6–1 (Cont.) Database Accounts

User Name	Description	See Also
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	The Oracle Text account.	<i>Oracle Text Reference</i>
DBSNMP	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	The account used by the Directory Integration Platform (DIP) to synchronize the changes in Oracle Internet Directory with the applications in the database.	None
DVSYST	There are two roles associated with this account. Database Vault owner role manages the Database Vault roles and configurations. The Database Vault Account Manager is used to manage database user accounts. Note: Part of Oracle Database Vault user interface text is stored in database tables in the DVSYST schema. By default, only the English language is loaded into these tables. You can use Oracle Database Vault Configuration Assistant to add more languages to Oracle Database Vault. For the necessary steps, refer to Appendix C in <i>Oracle Database Vault Administrator's Guide</i>	<i>Oracle Database Vault Administrator's Guide</i>
EXFSYST	The account owns the Expression Filter schema.	None
FLW_FILES	The account owns the Application Express uploaded files.	<i>Oracle Application Express Application Builder User's Guide</i>
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>
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	The Oracle Label Security administrator account.	<i>Oracle Label Security Administrator's Guide</i>
MDDATA	The schema used by Oracle Spatial for storing Geocoder and router data.	<i>Oracle Spatial Developer's Guide</i>
MDSYST	The Oracle Spatial and Oracle Multimedia Locator administrator account.	<i>Oracle Spatial Developer's Guide</i>
MGMT_VIEW	An account used by Oracle Enterprise Manager Database Control.	None

Table 6–1 (Cont.) Database Accounts

User Name	Description	See Also
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>
ORDPLUGINS	The Oracle Multimedia user. Plugins supplied by Oracle and third-party plugins are installed in this schema.	<i>Oracle Multimedia Reference</i>
ORDSYS	The Oracle Multimedia administrator account.	<i>Oracle Multimedia Reference</i>
ORDDATA	This account contains the Oracle Multimedia DICOM data model.	<i>Oracle Multimedia DICOM Developer's Guide</i>
OUTLN	The account that supports plan stability. Plan stability enables 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>
ORACLE_OCM	This account contains the instrumentation for configuration collection used by the Oracle Configuration Manager.	<i>Oracle Configuration Manager Installation and Administration Guide</i>
OWBSYS	The account used by Oracle Warehouse Builder as its default repository. You must unlock this account subsequent to installing the Oracle Database and prior to launching the Warehouse Builder Repository Assistant.	<i>Oracle Warehouse Builder Installation and Administration Guide for Windows and Linux</i>
OWBSYS_AUDIT	This account is used by the Warehouse Builder Control Center Agent to access the heterogeneous execution audit tables in the OWBSYS schema.	<i>Oracle Warehouse Builder Installation and Administration Guide for Windows and Linux</i>
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	An account used by Oracle sample programs and examples.	<i>Oracle Database Administrator's Guide</i>
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	The account that stores the information views for the SQL/MM Still Image Standard.	<i>Oracle Multimedia Reference</i>
SPATIAL_CSW_ADMIN_USR	The Catalog Services for the Web (CSW) account. It is used by the Oracle Spatial CSW cache manager to load all record type metadata, and record instances from the database into the main memory for the record types that are cached.	<i>Oracle Spatial Developer's Guide</i>

Table 6–1 (Cont.) Database Accounts

User Name	Description	See Also
SPATIAL_WFS_ADMIN_USR	The Web Feature Service (WFS) account. It is used by the Oracle Spatial WFS cache manager to load all feature type metadata, and feature instances from the database into main memory for the feature types that are cached.	<i>Oracle Spatial Developer's Guide</i>
SYS	The account used to perform database administration tasks.	<i>Oracle Database Administrator's Guide</i>
SYSMAN	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	Another account used to perform database administration tasks.	<i>Oracle Database Administrator's Guide</i>
WMSYS	The account used to store the metadata information for Oracle Workspace Manager.	<i>Oracle Database Workspace Manager Developer's Guide</i>
XDB	The account used for storing Oracle XML DB data and metadata.	<i>Oracle XML DB Developer's Guide</i>

Unlocking and Resetting 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, using one of the following methods:

- [Using Database Control to Unlock Accounts and Reset Passwords](#)
- [Using SQL*Plus to Unlock Accounts and Reset Passwords](#)
- [Unlocking and Changing Passwords](#)

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

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: The "[Logging In to Oracle Enterprise Manager Database Control](#)" section for information about logging in to the Database Control

2. Click **Server**.
3. In the Security section of the Server 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:

```
$ $ORACLE_HOME/bin/sqlplus
SQL> CONNECT SYS as SYSDBA
Enter password: SYS_password
```

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> PASSWORD account UNLOCK;
Changing password for account
New password: password
Retype new password: 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 the data through HTTP, unlock the ANONYMOUS user account.

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

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

Unlocking and Changing Passwords

Passwords for all Oracle system administration accounts except SYS, SYSTEM, SYSMAN, and DBSNMP are revoked after installation. Before you use a locked account, you must unlock it and reset its password. If you created a starter database during the installation, Oracle Database Configuration Assistant displays a screen with your database information and the Password Management button. Use the Password Management button to unlock only the user names you will use.

Apply the following guidelines when specifying passwords:

- Passwords must be between 8 and 30 characters long.
- Passwords must not start with a numeral.

- Passwords must not be the same as the user name.
- Passwords must not be Oracle reserved words.
- The SYS account password must not be `change_on_install`.
- The SYSTEM account password must not be `manager`.
- The SYSMAN account password must not be `sysman`.
- The DBSNMP account password must not be `dbsnmp`.
- If you choose to use the same password for all the accounts, then that password must not be `change_on_install`, `manager`, `sysman`, or `dbsnmp`.
- Passwords should have at least one alphabetic, one numeric, and one special character.
- Passwords should not be simple or obvious words, such as `welcome`, `account`, `database`, and `user`.
- Passwords should not have any consecutive repeating characters.

See Also: ["Reviewing Accounts and Passwords"](#) for more information about accounts and passwords

Identifying Databases

The Oracle Database 11g 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_world.example.com
```

In this example:

- `sales_world` is the name of the database. The database name portion is a string of no more than 30 characters that can contain alphanumeric, underscore (`_`), dollar (`$`), and pound (`#`) characters. The `DB_NAME` initialization parameter specifies the database name.
- `example.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.

Locating the Server 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 Automatic Storage Management for the database, Database Configuration Assistant 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: The "[Logging In to Oracle Enterprise Manager Database Control](#)" section for information about logging in to the Database Control

2. Click **Server**.
3. In the Database Configuration section of the Server page, click **Initialization Parameters**.

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

4. Select the **SPFile** tab.

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 on top of the page.

Reviewing Tablespaces and Data Files, Redo Log Files, and Control Files

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

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

Identifying Tablespaces and Data Files

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

Note: The `SYSAUX` and `SYSTEM` tablespaces must be present in all Oracle Database 11g databases.

[Table 6–2](#) describes the tablespaces provided by the default preconfigured database.

Table 6–2 *Tablespaces and Data Files*

Tablespace	Data File	Description
EXAMPLE	EXAMPLE01.DBF	Stores the Sample Schemas, if you included them.
SYSAUX	SYSAUX01.DBF	Serves as an auxiliary tablespace to the SYSTEM tablespace. Some products and options that previously used the SYSTEM tablespace now use the SYSAUX tablespace to reduce the load on the SYSTEM tablespace.
SYSTEM	SYSTEM01.DBF	Stores the data dictionary, including definitions of tables, views, and stored procedures needed by the Oracle Database. Information in this area is maintained automatically.
TEMP	TEMP01.DBF	Stores temporary tables and indexes created during the processing of your SQL statement. If you are running a SQL statement that involves a lot of sorting, such as the constructs GROUP BY, ORDER BY, or DISTINCT, then you may need to expand this tablespace.
UNDOTBS	UNDOTBS01.DBF	Stores undo information. The undo tablespace contains one or more undo segments that maintain transaction history that is used to roll back, or undo, changes to the database. All starter databases are configured to run in automatic undo management mode.
USERS	USERS01.DBF	Stores database objects created by database users.

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

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

1. Log in to the Database Control.

See Also: The "[Logging In to Oracle Enterprise Manager Database Control](#)" section for information about logging in to the Database Control

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

Enterprise Manager displays a table listing each data file, 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 11g 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: *Oracle Database Backup and Recovery User's Guide* for more information about redo log files

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

1. Log in to the Database Control.

See Also: The "[Logging In to Oracle Enterprise Manager Database Control](#)" section for information about logging in to the Database Control

2. Click **Server**.
3. In the Storage section of the Server 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 11g 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 data files and redo log files.

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

1. Log in to the Database Control.

See Also: "[Logging In to Oracle Enterprise Manager Database Control](#)" for information about logging in to the Database Control

2. Click **Server**.
3. In the Storage section of the Server page, click **Control Files**.

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 creating backups of these files to trace them, click **Help** in the Database Control window

For more information about setting the `CONTROL_FILES` initialization parameter value, refer to *Oracle Database Administrator's Guide*

Removing Oracle Database 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:

- [Reconfiguring Oracle Cluster Synchronization Services](#)
- [Removing Oracle Software Using the Deinstallation Tool](#)

See Also: *Oracle Grid Infrastructure Installation Guide* and *Oracle Real Application Clusters Installation Guide* for information on removing an Oracle RAC installation.

The "Dropping Disk Groups" section in the *Oracle Database Storage Administrator's Guide* for information on removing an Automatic Storage Management disk group.

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

Reconfiguring Oracle Cluster Synchronization Services

Oracle Cluster Synchronization Services (CSS) is a daemon process that is configured by the `root.sh` script when you configure a grid infrastructure instance. The CSS daemon runs out of the grid infrastructure home and is configured to start every time the system starts. This daemon process is required to enable synchronization between Oracle Automatic Storage Management and database instances. It must be running if an Oracle database is using Automatic Storage Management for database file storage.

Note: On cluster systems with Oracle RAC installations, the CSS daemon is configured during the Oracle Clusterware installation. If the system is running Oracle Clusterware, then refer to *Oracle Real Application Clusters Installation Guide* for information about removing Oracle RAC or Oracle Clusterware.

Removing Oracle Software Using the Deinstallation Tool

The `deinstall` command removes standalone Oracle Database installations, Oracle Clusterware and Automatic Storage Management (ASM) from your server, as well as Oracle Real Application Clusters (Oracle RAC) and Oracle Database client installations.

The following sections describe the command, and provide information about additional options to use the command:

- [About the Deinstallation Tool](#)
- [Example of Running the Deinstall Command for an Oracle Database](#)
- [Example of a Deinstallation Parameter File for Oracle Database](#)
- [Example of a Deinstallation Parameter File for Oracle Grid Infrastructure for a Standalone Server](#)
- [Removing Oracle Software](#)

Caution: If you have a standalone database on a node in a cluster and you have more than one database with the same global database name (GDN), then you cannot use the deinstall tool to remove one database only.

About the Deinstallation Tool

The Deinstallation Tool (`deinstall`) is available in the installation media before installation, and is available in Oracle home directories after installation. It is located in the path `$ORACLE_HOME/deinstall`.

The `deinstall` command stops Oracle software, and removes Oracle software and configuration files on the operating system.

The command uses the following syntax, where variable content is indicated in italics:

```
deinstall -home complete path of Oracle home [-silent] [-checkonly] [-local]
[-paramfile complete path of input parameter property file] [-params name1=value
name2=value . . .] [-o complete path of directory for saving files] [-help | -h]
```

The options are:

- `-silent`

Use this flag to run the command in silent or response file mode. If you use the `-silent` flag, then you must use the `-paramfile` flag, and provide a parameter file that contains the configuration values for the Oracle home that you want to deinstall or deconfigure.

You can generate a parameter file to use or modify by running `deinstall` with the `-checkonly` flag. The `deinstall` command then discovers information from the Oracle home that you want to deinstall and deconfigure. It generates the properties file, which you can then use with the `-silent` option.

You can also modify the template file `deinstall.rsp.tmpl`, located in the response folder.
- `-checkonly`

Use this flag to check the status of the Oracle software home configuration. Running the command with the `-checkonly` flag does not remove the Oracle configuration. The `-checkonly` flag generates a parameter file that you can use with the `deinstall` command.
- `-local`

Use this flag on a multinode environment to deinstall Oracle software in a cluster. When you run `deinstall` with this flag, it deconfigures and deinstalls the Oracle software on the local node (the node where `deinstall` is run). On remote nodes, it deconfigures Oracle software, but does not deinstall the Oracle software.

- `-paramfile` *complete path of input parameter property file*
Use this flag to run `deinstall` with a parameter file in a location other than the default. When you use this flag, provide the complete path where the parameter file is located.
The default location of the parameter file depends on the location of `deinstall`:
 - From the installation media or stage location: `$ORACLE_HOME/inventory/response`.
 - From a unzipped archive file from OTN: `/ziplocation/response`.
 - After installation from the installed Oracle home: `$ORACLE_HOME/deinstall/response`.
- `-params` [`name1=value name 2=value name3=value ...`]
Use this flag with a parameter file to override one or more values that you want to change in a parameter file you have already created.
- `-o` *complete path of directory for saving response files*
Use this flag to provide a path other than the default location where the properties file (`deinstall.rsp.tmpl`) is saved.
The default location of the parameter file depends on the location of `deinstall`:
 - From the installation media or stage location before installation: `$ORACLE_HOME/`
 - From a unzipped archive file from OTN: `/ziplocation/response/`.
 - After installation from the installed Oracle home: `$ORACLE_HOME/deinstall/response`.
- `-help` | `-h`
Use the help option (`-help` or `-h`) to obtain additional information about the command option flags.

Example of Running the Deinstall Command for an Oracle Database

As the `deinstall` command runs, you are prompted to provide the home directory of the Oracle software that you want to remove from your system. Provide additional information as prompted.

If you enter the `deinstall` command outside of the `$ORACLE_HOME/deinstall` folder, then help is displayed, unless you enter a `-home` flag and provide a path. If you run the `deinstall` command from the `$ORACLE_HOME/deinstall` folder, then deinstallation starts without prompting you for a home address.

Use the optional flag `-paramfile` to provide a path to a parameter file.

In the following example, the `deinstall` command is in the path `/u01/app/oracle/product/11.2.0/dbhome_1/deinstall`, and it uses a parameter file in the software owner location `/home/usr/oracle`:

```
# cd /u01/app/oracle/product/11.2.0/dbhome_1/deinstall
# ./deinstall -paramfile /home/usr/oracle/my_db_paramfile.tmpl
```

For the grid infrastructure home, use the `deinstall` script in the Oracle grid infrastructure for a standalone server home, which in this example is `/u01/app/oracle/product/11.2.0/grid`:

```
# cd /u01/app/oracle/product/11.2.0/grid/deinstall
```

```
# ./deinstall -paramfile /home/usr/oracle/my_grid_paramfile.tmpl
```

Example of a Deinstallation Parameter File for Oracle Database

You can run the `deinstall` command on a standalone Oracle Database with the `-paramfile` option to use the values you specify in the parameter file. The following is an example of a parameter file, in which the Oracle Database binary owner is `oracle`, the Oracle Database home (Oracle home) is in the path `/u01/app/oracle/product/11.2.0/dbhome_1/`, the Oracle base (where other Oracle software is installed) is `/u01/app/oracle/`, the central Oracle Inventory home (`oraInventory`) is `/u01/app/oraInventory`, the virtual IP address (VIP) is `192.0.2.1`, the local node (the node where you are running the deinstallation session from) is `myserver`, and the OSDBA group is `dba`:

```
#Copyright (c) 2005, 2006 Oracle Corporation. All rights reserved.
#Mon Feb 16 06:48:39 UTC 2009
DISK_GROUPS.sidb=
ASM_HOME=
ASM_LOCAL_SID=
LOGDIR=/u01/app/oracle/product/11.2.0/dbhome_1/oraInventory/logs/
ORACLE_BASE.sidb=/u01/app/oracle/
RECOVERY_LOC.sidb=
STORAGE_TYPE.sidb=FS
ORACLE_BASE=/u01/app/oracle/
INVENTORY_LOCATION=/u01/app/oraInventory
DB_TYPE.sidb=SI_DB
NODE_LIST.sidb=myserver
ARCHIVE_LOG_DESTINATION_LOC.sidb=
LOCAL_SID.sidb=sidb
DB_UNIQUE_NAME_LIST=sidb
ASM_FILES.sidb=
HOME_TYPE=SIDB
CRS_HOME=false
RAW_MAPPING_FILE.sidb=
SID_LIST.sidb=sidb
ORACLE_BINARY_OK=true
DATAFILE_LOC.sidb=/u01/app/oracle/oradata
local=false
LOCAL_NODE=myserver
CREATION_MODE.sidb=y
CONFIGFILE_LOC.sidb=
DIAG_DEST.sidb=/u01/app/oracle/
silent=false
ORACLE_HOME=/u01/app/oracle/product/11.2.0/dbhome_1/
SPFILE_LOC.sidb=
```

Example of a Deinstallation Parameter File for Oracle Grid Infrastructure for a Standalone Server

You can run the `deinstall` command on a Oracle grid infrastructure for a standalone server home with the `-paramfile` option to use the values you specify in the parameter file. The following is an example of a parameter file, in which the Oracle grid infrastructure binary owner is `oracle`, the Oracle grid infrastructure home is in the path `/u01/app/oracle/product/11.2.0/grid`, the Oracle base (where other Oracle software is installed) is `/u01/app/oracle/`, the central Oracle Inventory home (`oraInventory`) is `/u01/app/oraInventory`, the local node (the node where

you are running the deinstallation session from) is `myserver`, and the OSDBA group is `dba`:

```
#Copyright (c) 2005, 2009 Oracle Corporation. All rights reserved.
#Thu Mar 05 11:36:03 PST 2009
LOCAL_NODE=myserver
HOME_TYPE=SIHA
ASM_REDUNDANCY=EXTERNAL
ORACLE_BASE=/u01/app/oracle/
SCAN_PORT=0
silent=false
ASM_UPGRADE=false
ORA_CRS_HOME=/u01/app/oracle/product/11.2.0/grid
GPNPCONFIGDIR=$ORACLE_HOME
LOGDIR=/home/oracle/tmp/deinstall/logs/
ASM_DISCOVERY_STRING=/u02/stor/asm*
GPNPGCONFIGDIR=$ORACLE_HOME
ORACLE_OWNER=oracle
ASM_DISKSTRING=
CRS_STORAGE_OPTION=0
ORACLE_BINARY_OK=true
OCR_VOTINGDISK_IN_ASM=false
ASM_ORACLE_BASE=/u01/app/oracle
NETCFGJAR_NAME=netcfg.jar
ORA_DBA_GROUP=svrtech
JREDIR=/u01/app/oracle/grid/jdk/jre/
ORA_ASM_GROUP=dba
LANGUAGE_ID='AMERICAN_AMERICA.WE8ISO8859P1'
CSS_LEASEDURATION=400
ASM_HOME=/u01/app/oracle/grid
SHAREJAR_NAME=share.jar
HELPJAR_NAME=help4.jar
SILENT=false
local=false
INVENTORY_LOCATION=/u01/app/oraInventory
GNS_CONF=false
JEWTJAR_NAME=jewt4.jar
EMBASEJAR_NAME=oemlt.jar
ASM_
DISKS=/u02/stor/asm/asm0,/u02/stor/asm/asm2,/u02/stor/asm/asm3,/u02/stor/asm/asm1,
/u02/stor/asm/asm4,/u02/stor/asm/asm5,/u02/stor/asm/asm6,/u02/stor/asm/asm7,/u02/s
tor/asm/asm8
ORACLE_HOME=/u01/app/oracle/grid
CRS_HOME=true
ASM_IN_HOME=true
EWTJAR_NAME=ewt3.jar
ASM_DROP_DISKGROUPS=false
ASM_LOCAL_SID=+ASM
JLIBDIR=/u01/app/oracle/grid/jlib
VNDR_CLUSTER=false
ASM_DISK_GROUP=DATA
```

Removing Oracle Software

Complete the following procedure to remove Oracle software:

1. Log in as the installation owner.
2. Run the `deinstall` command, providing information about your servers as prompted.

Installing and Configuring Oracle Database Using Response Files

This appendix describes how to install and configure Oracle products using response files. It includes information about the following topics:

- [How Response Files Work?](#)
- [Creating the oraInst.loc File](#)
- [Preparing a Response File](#)
- [Running Oracle Universal Installer Using a Response File](#)
- [Running Net Configuration Assistant Using a Response File](#)
- [Running Database Configuration Assistant Using a Response File](#)
- [Postinstallation Configuration Using a Response File](#)

How Response Files Work?

You can automate the installation and configuration of Oracle software, either fully or partially, by specifying a response file when you start Oracle Universal Installer. Oracle Universal Installer uses the values contained in the response file to provide answers to some or all of Oracle Universal Installer prompt. It includes information about the following topics:

- [Reasons for Using Silent Mode or Response File Mode](#)
- [Creating a Database Using Automatic Storage Management as the Storage Option for Database Files](#)
- [General Procedure for Using Response Files](#)

Typically, Oracle Universal Installer runs in interactive mode, which means that it prompts you to provide information in graphical user interface (GUI) screens. When you use response files to provide this information, you run Oracle Universal Installer at a command prompt using either of the following modes:

- Silent mode

If you include responses for all of the prompts in the response file and specify the `-silent` option when starting Oracle Universal Installer, then Oracle Universal Installer runs in silent mode. During a silent-mode installation, Oracle Universal Installer does not display any screens. Instead, it displays progress information in the terminal that you used to start it.

- Response file mode

If you include responses for some or all of the prompts in the response file and omit the `-silent` option, then Oracle Universal Installer runs in response file mode. During a response file mode installation, Oracle Universal Installer displays all the screens, screens for which you specify information in the response file and also screens for which you did not specify the required information in the response file. The advantage is that you can validate the values in the screens for which you have already provided the information in the response file and continue with the installation.

You define the settings for a silent or response file installation by entering values for the variables listed in the response file. For instance, to specify the Oracle home name, you would supply the appropriate value for the `ORACLE_HOME` variable, as in the following example:

```
ORACLE_HOME="OraDBHome1"
```

Another way of specifying the response file's variable settings is to pass them as command line arguments when you run Oracle Universal Installer. For example:

```
-silent "ORACLE_HOME=OraDBHome1" ...
```

In this command, *directory_path* is the path of the database directory on the DVD or the path of the `Disk1` directory on the hard drive.

This method is particularly useful if you do not want to embed sensitive information, such as passwords, in the response file. For example:

```
-silent "s_dlgRBOPassword=binks342" ...
```

Ensure that you enclose the variable and its setting in quotes.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for more information about response file formats

Reasons for Using Silent Mode or Response File Mode

The following table describes several reasons why you might want to run Oracle Universal Installer in silent mode or response file 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>Oracle Universal Installer displays progress information in the terminal that you used to start it, but it does not display any of Oracle Universal Installer screens.</p>

Mode	Uses
Response File	<p>Use response file mode if you want to complete similar Oracle software installations on more than one system, providing default answers to some, but not all of Oracle Universal Installer prompts.</p> <p>In response file mode, all the installer screens are displayed, but defaults for the fields in these screens are provided by the response file. You have to provide information for the fields in screens where you have not provided values in the response file.</p>

Creating a Database Using Automatic Storage Management as the Storage Option for Database Files

Before you create a database that uses Automatic Storage Management, you must run the `root.sh` script. For this reason, you cannot create a database using Automatic Storage Management as the storage option for database files during a silent-mode installation. Instead, you can complete a software-only installation using silent mode, and then run the Oracle Net Configuration Assistant and Database Configuration Assistant 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 Automatic Storage Management as the storage option for database files. You can create a database that uses the file system option during a silent-mode installation.

General Procedure for Using Response Files

The following are the general steps to install and configure Oracle products using Oracle Universal Installer in silent or response file mode:

Note: You must complete all required preinstallation tasks on a system before running Oracle Universal Installer in silent or response file mode.

1. Create the `oraInst.loc` file.
2. Prepare a response file.
3. Run Oracle Universal Installer in silent or response file mode.
4. If you completed a software-only installation, then run Net Configuration Assistant and Database Configuration Assistant in silent or response file mode if required.

These steps are described in the following sections.

Creating the oraInst.loc File

If you plan to install Oracle products using Oracle Universal Installer in silent or response file mode, then 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 Oracle Universal 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 file.

To create the `oraInst.loc` file, follow these steps:

1. Switch user to root:

```
$ su - root
```

2. Create the `/etc/` directory if it does not exist:

```
# mkdir -p /var/opt/oracle
# mkdir /etc/
```

3. Change directory as follows:

```
# cd /etc/
```

4. Use a text editor to create the `oraInst.loc` file, containing the following lines:

```
inventory_loc=$ORACLE_BASE/oraInventory
inst_group=oinstall
```

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
```

Preparing a Response File

This section describes the following methods to prepare a response file for use during silent mode or response file mode installations:

- [Editing a Response File Template](#)
- [Recording a Response File](#)

Editing a Response File Template

This method is most useful for the Enterprise Edition or Standard Edition installation types.

Oracle provides response file templates for each product and installation type, and for each configuration tool. These files are located at `database/response` directory on the installation media.

Note: If you copied the software to a hard disk, the response files are located in the `Disk1/response` directory.

[Table A-1](#) lists the response files provided with Oracle Database.

Table A-1 Response Files

Response File	Description
db_install.rsp	Response file installation of Oracle Database 11g
dbca.rsp	Response file installation of Database Configuration Assistant
netca.rsp	Response file installation of Oracle Net Configuration Assistant

To copy and modify 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 path to the database directory on the installation media. If you have copied the software to a hard drive, then 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
```

Remember that you can specify sensitive information, such as passwords, at the command line rather than within the response file. "How Response Files Work?" explains this method.

See Also: *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for detailed information on creating response files

3. Follow the instructions in the file to edit it.

Note: Oracle Universal Installer or configuration assistant fails if you do not correctly configure the response file. Refer to "Silent-Mode Response File Error Handling" section on page G-5 for more information about troubleshooting a failed response file mode installation.

4. Change the permissions on the file to 700:

```
$ chmod 700 /local_dir/response_file.rsp
```

Note: A fully specified response file for an Oracle Database installation contains the passwords for database administrative accounts and for a user who is a member of the OSDBA group (required for automated backups). Ensure 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 Oracle Universal Installer in interactive mode to record a response file, which you can edit and then use to complete silent mode or response file mode installations. This method is useful for custom or software-only installations.

Starting with Oracle Database 11g Release 2 (11.2), you can save all the installation steps into a response file during installation. You can click the Save Response File button on the Summary page to do this. Later, this file can be used for a silent installation.

When you record the response file, you can either complete the installation, or you can exit from Oracle Universal Installer on the Summary page, before it starts to copy the software to the system.

If you use record mode during a response file mode installation, then Oracle Universal Installer records the variable values that were specified in the original source response file into the new response file.

Note: You cannot use record mode to create a response file during an installation that uses the Typical installation method.

To record a response file:

1. Complete the preinstallation tasks listed in [Chapter 2](#).

When you run Oracle Universal Installer to record a response file, it checks the system to verify that it meets the requirements to install the software. For this reason, Oracle recommends that you complete all of the required preinstallation tasks and record the response file while completing an installation.

2. If you have not installed Oracle software on this system previously, create the `oraInst.loc` file as described in [Creating the oraInst.loc File](#).
3. Ensure that the Oracle software owner user has permissions to create or write to the Oracle home path that you will specify when you run Oracle Universal Installer.
4. On each Oracle Universal Installer screen, specify the required information.

See Also: [Running Oracle Universal Installer](#) for information on the installation process

5. When Oracle Universal Installer displays the Summary screen, perform the following:
 1. Click **Save Response File** and specify a file name and location to save the values for the response file.
 2. Click **Finish** to create the response file and continue with the installation.
 3. Click **Cancel** if you only want to create the response file but not continue with the installation. The installation will stop, but the settings you have entered will be recorded in the response file.
6. If you do not complete the installation, then delete the Oracle home directory that Oracle Universal Installer created using the path you specified in the Specify File Locations screen.
7. Before you use the saved response file on another system, edit the file and make any required changes.

Use the instructions in the file as a guide when editing it.

Running Oracle Universal Installer Using a Response File

Now, you are ready to run Oracle Universal Installer at the command line, specifying the response file you created, to perform the installation. The Oracle Universal Installer executable, `runInstaller`, provides several options. For help information on the full set of these options, run the `runInstaller` command with the `-help` option, for example:

```
$ directory_path/runInstaller -help
```

The help information appears in a window after some time.

To run Oracle Universal Installer using a response file:

1. Complete the preinstallation tasks listed in [Chapter 2](#).
2. Log in as the Oracle software owner user (typically, `oracle`).
3. If you are completing a response file 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 Oracle Universal Installer in silent or response file mode, enter a command similar to the following:

```
$ directory_path/runInstaller [-silent] [-noconfig] \  
-responseFile responsefilename
```

Note: Do not specify a relative path to the response file. If you specify a relative path, then Oracle Universal Installer fails.

In this example:

- *directory_path* is the path of the database directory on the DVD or the path of the `Disk1` directory on the hard drive.
- `-silent` indicates that you want to run Oracle Universal Installer in silent mode.
- `-noconfig` suppresses running the configuration assistants during installation, and a software-only installation is performed instead.
- *responsefilename* is the full path and file name of the installation response file that you configured.

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:

```
$ sudo sh  
password:
```

```
# /oracle_home_path/root.sh
```

Running Net Configuration Assistant Using a Response File

You can run Net Configuration Assistant in silent mode to configure and start an Oracle Net listener on the system, configure naming methods, and configure Oracle Net service names. To run Net Configuration Assistant 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 in the `database/response` directory on the DVD.

Note: If you copied the software to a hard disk, then the response file template is located in the `Disk1/response` directory.

To run Net Configuration Assistant using a response file:

1. Copy the `netca.rsp` response file template from the response file directory to a directory on your system:

```
$ cp /directory_path/response/netca.rsp local_directory
```

In this example, `directory_path` is the path of the `database` directory on the DVD. 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. Follow the instructions in the file to edit it.

Note: Net Configuration Assistant fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.
5. Enter a command similar to the following to run Net Configuration Assistant in silent mode:

```
$ $ORACLE_HOME/bin/netca /silent /responsefile /local_dir/netca.rsp
```

In this command:

- The `/silent` option indicates that you want to run Net Configuration Assistant in silent mode.
- `local_dir` is the full path of the directory where you copied the `netca.rsp` response file template.

Running Database Configuration Assistant Using a Response File

You can run Database Configuration Assistant in response file or silent mode to configure and start an Oracle Database on the system. To run Database Configuration Assistant in response file or silent mode, you must copy and edit a response file template. Oracle provides a response file template named `dbca.rsp` in the `database/response` directory on the DVD.

Note: If you copied the software to a hard disk, then the response file template is located in the `Disk1/response` directory.

This section contains the following topics:

- [Using Database Configuration Assistant in Response File Mode](#)
- [Using Database Configuration Assistant in Silent Mode](#)
- [Running Database Configuration Assistant in Response File or Silent Mode](#)

Using Database Configuration Assistant in Response File Mode

Use `-progressOnly` flag to set the mode to response file. In the response file mode, Database Configuration Assistant 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 Database Configuration Assistant in response file mode, you must use a graphical display and set the `DISPLAY` environment variable.

Using Database Configuration Assistant in Silent Mode

Use `-silent` flag to set the mode to silent. In the silent mode, Database Configuration Assistant uses values that you specify, in the response file or as command line options, to create a database.

Running Database Configuration Assistant in Response File or Silent Mode

To run Database Configuration Assistant in response file or silent mode:

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 Database Configuration Assistant. 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 path of the database directory on the DVD. 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: Database Configuration Assistant fails if you do not correctly configure the response file.

4. Log in as the Oracle software owner user, and set the `ORACLE_HOME` environment variable to specify the correct Oracle home directory.
5. If you intend running Database Configuration Assistant in response file mode, set the `DISPLAY` environment variable.
6. Enter a command similar to the following to run Database Configuration Assistant in response file or silent mode with a response file:

```
$ORACLE_HOME/bin/dbca {-progressOnly | -silent} -responseFile \  
/local_dir/dbca.rsp
```

In this example:

- The `-silent` option indicates that you want to run Database Configuration Assistant in silent mode.
- The `-progressOnly` option indicates that you want to run Database Configuration Assistant in response file mode.
- `local_dir` is the full path of the directory where you copied the `dbca.rsp` response file template.

Postinstallation Configuration Using a Response File

Use the following sections to create and run a response file configuration after installing Oracle software.

About the Postinstallation Configuration File

When you run a silent or response file installation, you provide information about your servers in a response file that you otherwise provide manually during a graphical user interface installation. However, the response file does not contain passwords for user accounts that configuration assistants require after software installation is complete. The configuration assistants are started with a script called `configToolAllCommands`. You can run this script in response file mode by using a password response file. The script uses the passwords to run the configuration tools in succession to complete configuration.

If you keep the password file to use for clone installations, then Oracle strongly recommends that you store it in a secure location. In addition, if you have to stop an installation to fix an error, you can run the configuration assistants using `configToolAllCommands` and a password response file.

The `configToolAllCommands` password response file consists of the following syntax options:

- `internal_component_name` is the name of the component that the configuration assistant configures
- `variable_name` is the name of the configuration file variable
- `value` is the desired value to use for configuration.

The command syntax is as follows:

```
internal_component_name | variable_name=value
```


For example:

```
oracle.assistants.asm|S_ASMPASSWORD=welcome
```

Oracle strongly recommends that you maintain security with a password response file:

- Permissions on the response file should be set to 600.
- The owner of the the response file should be the installation owner user, with the group set to the central inventory (oraInventory) group.

Running Postinstallation Configuration Using a Response File

To run configuration assistants with the `configToolAllCommands` script:

1. Create a response file using the syntax `filename.properties`. For example:

```
$ touch cfgrsp.properties
```

2. Open the file with a text editor, and cut and paste the password template, modifying as needed.

Example A–1 Password response file for Oracle grid infrastructure for a standalone server

Oracle grid infrastructure requires passwords for Automatic Storage Management Configuration Assistant (ASMCA), and for Intelligent Platform Management Interface Configuration Assistant (IPMICA) if you have a BMC card and you want to enable this feature. Provide the following response file,

```
oracle.assistants.asm|S_ASMPASSWORD=password
oracle.assistants.asm|S_ASMMONITORPASSWORD=password
```

Example A–2 Password response file for Oracle Database

Oracle Database configuration requires the SYS, SYSTEM, SYSMAN, and DBSNMP passwords for use with Database Configuration Assistant (DBCA). The `S_ASMSNMPPASSWORD` response is necessary only if the database is using ASM for storage. Also, if you selected to configure Oracle Enterprise Manager, then you must provide the password for the Oracle software installation owner for the `S_HOSTUSERPASSWORD` response.

```
oracle.assistants.server|S_SYSPASSWORD=password
oracle.assistants.server|S_SYSTEMPASSWORD=password
oracle.assistants.server|S_SYSMANPASSWORD=password
oracle.assistants.server|S_DBSNMPPASSWORD=password
oracle.assistants.server|S_HOSTUSERPASSWORD=password
oracle.assistants.server|S_ASMSNMPPASSWORD=password
```

If you do not want to enable Oracle Enterprise Manager or ASM, then leave those password fields blank

3. Change permissions to secure the file. For example:

```
$ ls -al cfgrsp
-rw----- 1 oracle oinstall 0 Apr 30 17:30 cfgrsp
```

4. Change directory to `$ORACLE_HOME/cfgtoollogs`

Run the configuration script using the following syntax:

```
configToolAllCommands RESPONSE_FILE=/path/name.properties
```

for example:

```
$ ./configToolAllCommands RESPONSE_FILE=/home/oracle/cfgrsp.properties
```

Cloning an Oracle Home

Cloning an Oracle home involves creating a copy of the Oracle home and then configuring it for a new environment. If you are performing multiple Oracle Database installations, then you may want to use this method to create each Oracle home, because copying files from an existing Oracle Database installation takes less time than creating a new version of them. This method is also useful if the Oracle home that you are cloning has had patches applied to it. When you clone this Oracle home, the new Oracle home will have the patch updates as well.

When cloning Database Oracle homes using 11.2 Database Control, you must update the exclude file list. This file list specifies files that need not be included when the source Oracle home is archived because these files are not required for the clone operation. The following files should not be included in the archive:

- `sqlnet.ora`
- `tnsnames.ora`
- `listener.ora`
- `oratab`

Note: In addition to cloning an Oracle home, you can clone individual Oracle Database installations by using Enterprise Manager Database Control. *Oracle Database Administrator's Guide* provides detailed information about cloning Oracle Database installations and Oracle homes.

This appendix includes information about the following topics:

- [Cloning an Oracle Home](#)
- [Configuring Oracle Configuration Manager in a Cloned Oracle Home](#)

Cloning an Oracle Home

Perform the following to clone an Oracle home:

1. Verify that the installation of Oracle Database that you want to clone has been successful.

You can do this by reviewing the `installActionsdate_time.log` file for the installation session, which is normally located in the `/orainventory_location/logs` directory.

If you have installed patches, then you can check their status by running the following commands:

```
$ $ORACLE_HOME/OPatch ORACLE_HOME=ORACLE_HOME_using_patch  
$ $ORACLE_HOME/OPatch opatch lsinventory
```

2. Stop all processes related to the Oracle home. Refer to "[Removing Oracle Software](#)" for more information on stopping the processes for an Oracle home.
3. Create a ZIP file with the Oracle home (but not Oracle base) directory.

For example, if the source Oracle installation is in the `/u01/app/oracle/product/11.2.0/dbhome_1`, then you zip the `dbhome_1` directory by using the following command:

```
# zip -r dbhome_1.zip /u01/app/oracle/product/11.2.0/dbhome_1
```

Do not include the `admin`, `recovery_area`, and `oradata` directories that are under the Oracle base directory. These directories will be created in the target installation later, when you create a new database there.

4. Copy the ZIP file to the root directory of the target computer.
5. Extract the ZIP file contents by using the following command:

```
# unzip -d / dbhome_1.zip
```
6. Repeat steps 4 and 5 for each computer where you want to clone the Oracle home, unless the Oracle home is on a shared storage device.
7. On the target computer, change directory to the unzipped Oracle home directory, and remove all the `.ora` (`*.ora`) files present in the unzipped `$ORACLE_HOME/network/admin` directory.
8. From the `$ORACLE_HOME/clone/bin` directory, run `clone.pl` for the unzipped Oracle home. Use the following syntax:

```
$ORACLE_HOME/perl/bin/perl clone.pl ORACLE_BASE="target_oracle_base" ORACLE_  
HOME="target_oracle_home"  
OSDBA_GROUP=OSDBA_privileged_group OSOPER_GROUP=OSOPER_privileged_group  
-defaultHomeName
```

For example:

```
$ORACLE_HOME/perl/bin/perl clone.pl ORACLE_BASE="/u01/app/oracle/" ORACLE_  
HOME="/u01/app/oracle/product/11.2.0/dbhome_1"  
OSDBA_GROUP=dba OSOPER_GROUP=oper -defaultHomeName
```

Oracle Universal Installer starts, and then records the cloning actions in the `cloneActionstimestamp.log` file. This log file is normally located in `/orainventory_location/logs` directory.

9. To configure connection information for the new database, run Net Configuration Assistant:

```
$ cd $ORACLE_HOME/bin  
$ ./netca
```

10. To create a new database for the newly cloned Oracle home, run Database Configuration Assistant as follows:

```
$ cd $ORACLE_HOME/bin  
$ ./dbca
```

See Also:

- *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for detailed information about using Oracle Universal Installer to clone an Oracle Database home
- *Oracle Database Administrator's Guide* for information about cloning an Oracle databases, as well as cloning an Oracle Database home

Configuring Oracle Configuration Manager in a Cloned Oracle Home

Perform the following to configure Oracle Configuration Manager for a cloned Oracle home:

1. Run the following command from `$ORACLE_HOME/ccr/state`:

```
$ rm -rf *.ll*
```

2. Run the following command from `$ORACLE_HOME/ccr/inventory`:

```
$ cp core.jar pending
```

3. Run the following commands from `$ORACLE_HOME/ccr/bin`:

```
$ ./emSnapshotEnv
```

```
$ ./deployPackages
```

4. Run the following command from `$ORACLE_HOME/ccr/bin`, and provide the proper credentials:

```
$ ./configCCR
```

If Oracle Configuration Manager was manually configured using `setupCCR`, then perform the following in the cloned Oracle home:

1. Delete all the subdirectories of the `$ORACLE_HOME/ccr/hosts` directory to remove previously configured hosts.
2. Run the following command from `$ORACLE_HOME/ccr/bin`:

```
$ ./configCCR -a
```

If you have installed Oracle Configuration Manager in the original Oracle home but have not configured it, then run the following command in the cloned Oracle home:

```
$ setupCCR
```

Using NAS Devices

This appendix provides guidelines for using a network attached storage (NAS) storage device for Oracle software and database files. It includes information about the following:

- [General Configuration Guidelines for NAS Devices](#)
- [NFS Feature Description](#)
- [Choosing Mount Points](#)
- [Creating Files on a NAS Device for Use with Automatic Storage Management](#)
- [NFS Mount Options](#)

See Also: ["Configuring and Using Direct NFS Client"](#)

General Configuration Guidelines for NAS Devices

Refer to the documentation provided with the 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 the requirements:

- Before using the NAS device for the installation, verify that it is certified.

Note: NAS devices are certified only for single-instance databases. For more information on this refer to note 359515.1 on the My Oracle Support (formerly Oracle*MetaLink*) Web site:

<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 instance database installations, Oracle recommends that you 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.

NFS Feature Description

The following are the features of NFS:

- Oracle kernel handles best possible configuration to perform optimal I/O using available resources. This enables better configuration management.
- NFS storage is now available across different platforms like Windows.
- ODM NFS helps standardize all the tunable configuration parameters.
- ODM NFS has a stable NFS client that does not affect kernel performance. It optimizes the I/O path when making NFS operations. This ensures higher stability.
- NFA provides better diagnostics in case of errors.

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 the following sections comply with the Optimal Flexible Architecture recommendations:

- [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 (<code>\$ORACLE_BASE</code>)	<code>/u01/app/oracle</code>
Oracle Inventory	<code>\$ORACLE_BASE/oraInventory</code>
Oracle home	<code>\$ORACLE_BASE/product/11.2.0/dbhome_1</code>

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, then 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, then 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 Oracle Universal 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, then 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/recovery_area`)

Oracle recommends that you use different file systems for database and recovery files.

If you use this mount point, then 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/10.2.0/dbhome_1
/u01/app/oracle/product/11.2.0/dbhome_1
```

- The release directory (`oracle_base/product/11.2.0`)

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/11.2.0/dbhome_1
/u01/app/oracle/product/11.2.0/client_1
```

- The Oracle home directory (`oracle_base/product/11.2.0/dbhome_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/11.2.0/dbhome_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, then use paths or mount points similar to the following:

File Type	Path or Mount Point
Database files	<code>/u02/oradata</code>
Recovery files	<code>/u03/recovery_area</code>

When Oracle Universal Installer prompts you for the data file and the recovery file directories, specify these paths. The Database Configuration Assistant 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/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, then 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/recovery_area/orcl
```

Specify the directory `/u02/oradata` when Oracle Universal Installer prompts you for the data file directory and specify the directory `/u03/recovery_area` when Oracle Universal Installer prompts you for the recovery file location. The `orcl` directory will be used automatically either by Database Configuration Assistant or by Enterprise Manager.

Creating Files on a NAS Device for Use with Automatic Storage Management

If you have a certified NAS storage device, then you can create zero-padded files in an NFS mounted directory and use those files as disk devices in an Automatic Storage Management disk group. To create these files, follow these steps:

Note: To use files as disk devices in an Automatic Storage Management 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.

Refer to the NAS device documentation for more information about completing this step.

2. Switch user to `root`:

```
$ sudo sh
password:
```

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 restarts, add an entry for the file system in the `/etc/mtab` mount file.

For more information about editing the mount file for the operating system, refer to the man pages. For more information about recommended mount options, refer to the "NFS Mount Options" section on page C-5.

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 to create the required number of zero-padded files in this directory:

```
# dd if=/dev/zero of=/mnt/oracleasm/nfsdg/disk1 bs=1024k count=1000
```

This example creates 1 GB files on the NFS file system. You must create one, two, or three files respectively to create an external, normal, or high redundancy disk group.

Note: Creating multiple zero-padded files on the same NAS box does not guard against NAS box failure. Instead, create one file for each NAS box and mirror them using the Automatic Storage Management technology.

9. 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
```

10. When you are creating the database, edit the Automatic Storage Management 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

You must mount NFS volumes used for storing database files with special mount options on the host where the database server is running. When mounting an NFS file system, Oracle recommends that you use the same mount point options that the NAS

vendor used when certifying the device. Refer to the device documentation or contact the vendor for information about recommended mount-point options.

Option	Requirement	Description
hard	Mandatory	Generate a hard mount of the NFS file system. If the connection to the server fails or is temporarily lost, then connection attempts are made until the NAS device responds.
bg	Optional	Try to connect in the background if connection fails.
tcp	Optional	Use the TCP protocol rather than UDP. TCP is more reliable than UDP.
nfsvers=3	Optional	Use NFS version 3. Oracle recommends that you use NFS version 3 where available, unless the performance of version 2 is later.
suid	Optional	Allow clients to run executables with SUID enabled. This option is required for Oracle software mount points.
rsize	Mandatory	The number of bytes used when reading from the NAS device. This value should be set to the maximum database block size supported by this platform. A value of 8192 is often recommended for NFS version 2 and 32768 is often recommended for NFS version 3.
wsize	Mandatory	The number of bytes used when writing to the NAS device. This value should be set to the maximum database block size supported by this platform. A value of 8192 is often recommended for NFS version 2 and 32768 is often recommended for NFS version 3.
nointr (or intr)	Optional	Do not allow (or allow) keyboard interrupts to end 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 the vendor for advice.
actime=0 or noac	Mandatory	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, then Oracle Universal Installer will not install the software in the directory that you specify.
timeo	Mandatory	Represents number of seconds to timeout. Better overall performance may be achieved by increasing the timeout when mounting on a busy network, to a slow server, or through several routers or gateways. Oracle recommends that you set the timeout value to the maximum timeout for TCP, which is 600 seconds.
rw	Mandatory	Represents read and write access.

Optimal Flexible Architecture

This appendix describes the Optimal Flexible Architecture standard. The standard is a set of configuration guidelines created to ensure well organized Oracle installations that are easier to maintain. It includes information about the following topics:

- [Overview of the Optimal Flexible Architecture Standard](#)
- [Implementing Optimal Flexible Architecture](#)

Overview of the Optimal Flexible Architecture Standard

All Oracle components on the installation media are compliant with Optimal Flexible Architecture. This means, Oracle Universal Installer places Oracle Database components in directory locations that follow Optimal Flexible Architecture guidelines.

Although using Optimal Flexible Architecture is not a requirement, Oracle recommends that you use it if the database will grow in size, or if you plan to have multiple databases.

Implementing Optimal Flexible Architecture

This section describes the naming strategy recommended by the Optimal Flexible Architecture standard. It contains the following sections:

- [File Systems](#)
- [Naming Directories](#)
- [Naming Database Files](#)
- [Separating Segments with Different Requirements](#)
- [Exploiting the Optimal Flexible Architecture Structure for Oracle Files](#)
- [Optimal Flexible Architecture File Mapping](#)

File Systems

The following sections describe the conventions for mount points:

- [Number of File Systems](#)
- [Naming Conventions](#)

Number of File Systems

To fully implement the Optimal Flexible Architecture 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 Directories

The following sections describe the naming conventions for directories that are compliant with the Optimal Flexible Architecture standard:

- [Oracle Base Directory Naming Convention](#)
- [Naming Mount Points for Very Large Databases \(VLDBs\)](#)
- [Referring to Path Names](#)
- [Oracle Home Directory Naming Convention](#)
- [Naming Subdirectories](#)

Note: Ensure that the paths you select for Oracle software, such as the Oracle home path and the Oracle base path, use only ASCII characters. Because installation owner names are used by default for some path, this ASCII character restriction applies to user names, file names, and directory names.

Oracle Base Directory Naming Convention

The Oracle Base directory is the top level directory that you can use to install the various Oracle software products. You can use the same Oracle base directory for more than one installation. If different operating system users install Oracle software on the same system, then each user must create a separate Oracle base directory.

Name Oracle base directories using the syntax `/pm/s/u`. [Table D-1](#) describes the variables used in this syntax.

Table D-1 *Syntax for Naming Oracle Base Directories*

Variable	Description
<code>pm</code>	A mount point name
<code>s</code>	A standard directory name
<code>u</code>	The name of the owner of the directory (the user running Oracle Universal Installer)

For example, `/u01/app/oracle` is an Oracle base directory created by the `oracle` user and `/u01/app/applmgr` is an Oracle base directory created by the `applmgr` user.

Placing Oracle base directories at the same level in the UNIX file system is advantageous because it enables you to refer to the collection of Oracle base directories on different mount points using a single pattern matching string, `/*/app/*`.

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 `/h/q/d` for naming mount points. [Table D-2](#) describes the variables used in this syntax.

Table D-2 Syntax for Naming Mount Points for Very Large Databases

Variable	Description
<i>h</i>	Oracle base directory
<i>q</i>	A string denoting that Oracle data is stored in this directory, for example, <code>oradata</code>
<i>d</i>	The value of the initialization parameter <code>DB_NAME</code> (typically the same as the instance <code>SID</code> for single-instance databases)

For example, to allocate two drives exclusively for the `test` database, name the mount points `/u01/app/oracle/oradata/test` and `/u02/app/oracle/oradata/test`.

Referring to Path Names

Refer to explicit path names 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 Optimal Flexible Architecture requirement of simultaneously running multiple versions of Oracle software, install the software in a directory matching the pattern `/pm/s/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>s</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 (<code>dbhome_1</code>), Client (<code>client</code>), or Oracle grid infrastructure (<code>grid</code>)
<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/11.2.0/dbhome_1` indicates the Oracle home directory for the first installation of Oracle Database on this system.

The `ORACLE_HOME` environment variable is set to 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
arch	Archived redo log files
adump	Audit files (Set the <code>AUDIT_FILE_DEST</code> initialization parameter to specify the adump directory. Clean out this subdirectory periodically)
create	Contains the data pump file <code>dp.log</code>
dpdump	Default directory for data pump operations. Scripts used to create the database
exp	Database export files
logbook	Files recording the status and history of the database
pfile	Instance parameter files
scripts	Ad hoc SQL scripts

For example, `/u01/app/oracle/admin/orcl/scripts/` is the `scripts` subdirectory associated with the database named `orcl`.

In Oracle Database 11g, Automatic Diagnostic Repository (ADR) directories replace the `bdump`, `cdump`, and `udump` directories. The ADR diagnostic data will go into the `/h/diag/rdbms/d/i/` directory.

where

`h` is Oracle Base

`d` is the database name

`i` is the instance name.

Under ADR home we have the trace, alert, and incident sub-directories. [Table D-5](#) describes the ADR directories.

Table D-5 Locations for Diagnostic Traces

Diagnostic Data	10g Location	11g Location
Foreground Process traces	<code>user_dump_dest</code>	<code>ADR_HOME/trace/</code>
Background Process traces	<code>background_dump_dest</code>	<code>ADR_HOME/trace/</code>
Alert Log Data	<code>background_dump_dest</code>	<code>ADR_HOME/alert/</code>
Core Dump	<code>core_dump_dest</code>	<code>ADR_HOME/incident/In/</code>
Incident Dumps	<code>user_dump_dest</code> or <code>background_dump_dest</code> depending on the process	<code>ADR_HOME/incident/In/</code>

Naming Database Files

The following table lists the recommended file naming conventions for database files:

Note: Oracle Managed Files (OMF) and files stored in Automatic Storage Management 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>/h/q/d/control.ctl</code>
Redo log files	<code>/h/q/d/redon.log</code>
Data files	<code>/h/q/d/tn.dbf</code>

The following table describes this syntax:

Variable	Description
<i>h</i>	Oracle base directory
<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 data files associated with database *d* in the path `/h/q/d`.

Using this convention, it is easy to determine the database to which the `/u01/app/oracle/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-6](#) 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: *Oracle Database Administrator's Guide* for information about creating databases manually

Table D-6 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

Table D-6 (Cont.) Special Tablespaces

Tablespace	Required	Description
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.

See Also: ["Reviewing Tablespaces and Data Files, Redo Log Files, and Control Files"](#) for information on redo log, and control files

Exploiting the Optimal Flexible Architecture Structure for Oracle Files

[Table D-7](#) describes the syntax used for identifying classes of files.

Table D-7 Directory Structure Syntax for Identifying Classes of Files

Directory Structure Syntax	Description
<code>/u[0-9][0-9]</code>	User data directories
<code>*/home/*</code>	User home directories
<code>*/app/*</code>	User application software directories
<code>*/app/applmgr</code>	Oracle applications software subtrees
<code>*/app/oracle/product</code>	Oracle software subtrees
<code>*/app/oracle/product/11.2.0</code>	Oracle software subtree for release 11g products
<code>*/app/oracle/product/11.2.0/db*</code>	Oracle home directories for Oracle Database 11g
<code>*/app/oracle/product/11.2.0/grid*</code>	Oracle home directory for Oracle grid infrastructure 11g for user oracle
<code>*/app/oracle/admin/orcl</code>	orcl database administrative subtrees
<code>*/app/oracle/admin/orcl/arch/*</code>	orcl database archived log files
<code>*/app/oracle/oradata</code>	Oracle data directories
<code>*/app/oracle/oradata/orcl/*</code>	orcl database files
<code>*/app/oracle/oradata/orcl/*.log</code>	orcl database redo log files

Optimal Flexible Architecture File Mapping

[Table D-8](#) shows a hierarchical file mapping of a sample Optimal Flexible Architecture-compliant installation with two Oracle home directories and two databases. The database files are distributed across three mount points, `/u02`, `/u03`, and `/u04`.

Note: Oracle recommends that you use ASM to provide greater redundancy and throughput.

Table D–8 Hierarchical File Mapping for an Optimal Flexible Architecture Installation

Directory	Description
/	Root directory
/u01/	User data mount point 1
/u01/app/	Subtree for application software
/u01/app/oracle/	Oracle Base directory
/u01/app/oracle/admin/	Subtree for database administration files
/u01/app/oracle/admin/TAR	Subtree for support log files
/u01/app/oracle/admin/db_name1/	admin subtree for <i>db_name1</i> database
/u01/app/oracle/admin/db_name2/	admin subtree for <i>db_name2</i> database
/u01/app/oracle/doc/	Online documentation
/u01/app/oracle/recovery_area/	Subtree for recovery files
/u01/app/oracle/recovery_area/db_name1	Recovery files for <i>db_name1</i> database
/u01/app/oracle/recovery_area/db_name2	Recovery files for <i>db_name2</i> database
/u02/app/oracle/oradata	
/u03/app/oracle/oradata	Oracle data directory
/u04/app/oracle/oradata	
/u01/app/oracle/product/	Distribution files
/u01/app/oracle/product/11.2.0/dbhome_1	Oracle home directory for Oracle Database, for user <code>oracle</code>
/u01/app/oracle/product/11.2.0/grid	Oracle home directory for Oracle grid infrastructure, for user <code>oracle</code>
/u01/app/kjff/	Oracle base directory for user <code>kjff</code>
/u01/app/edm/	Oracle base directory for user <code>edm</code>

Managing Oracle Database Port Numbers

During installation, Oracle Universal 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:

- [About Managing Ports](#)
- [Viewing Port Numbers and Access URLs](#)
- [Port Numbers and Protocols of Oracle Components](#)
- [Changing the Oracle Enterprise Management Agent Port](#)
- [Changing the Oracle Enterprise Manager Database Control Ports](#)
- [Changing the Oracle XML DB Ports](#)

About Managing Ports

During installation, Oracle Universal Installer assigns port numbers to components from a set of default port numbers. Many Oracle Database components and services use ports. As an administrator, it is important to know the port numbers used by these services, and to make sure that the same port number is not used by two services on your host. Enter the following command to identify the ports currently used on your computer:

```
$/bin/netstat -a
```

Most port numbers are assigned during installation. Every component and service has an allotted port range, which is the set of port numbers Oracle Database attempts to use when assigning a port. Oracle Database starts with the lowest number in the range and performs the following checks:

- Is the port used by another Oracle Database installation on the host?
The installation may be up or down at the time; Oracle Database can still detect if the port is used.
- Is the port used by a process that is currently running?
This could be any process on the host, even a non-Oracle Database process.
- Is the port listed in the `/etc/services` files?

If the answer to any of the preceding questions is *yes*, Oracle Database moves to the next highest port in the allotted port range and continues checking until it finds a free port.

Viewing Port Numbers and Access URLs

In most cases, the Oracle Database component's port number is listed in the tool used to configure the port. In addition, ports for some Oracle Database applications are listed in the `portlist.ini` file. This file is located in the `$ORACLE_HOME/install` directory.

If you change a port number, it is not updated in the `portlist.ini` file, so you can only rely on this file immediately after installation. To find or change a port number, use the methods described in this appendix.

Port Numbers and Protocols of Oracle Components

The following table lists the port numbers and protocols 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.

Table E-1 Ports Used in Oracle Components

Component and Description	Default Port Number	Port Range	Protocol
Oracle SQL*Net Listener/Oracle Net Enables Oracle client connections to the database over Oracle's SQL*Net protocol. You can configure it during installation. To reconfigure this port, use Net Configuration Assistant.	1521	1024-65535	TCP
Connection Manager Listening port for Oracle client connections to Oracle Connection Manager. It is not configured during installation, but can be configured manually by editing the <code>cman.ora</code> parameter file. You can find the file under <code>/network/admin</code> directory.	1630	1630	TCP
Oracle Enterprise Manager Database Control HTTP port for Enterprise Manager Database Control. It is configured during installation. "Changing the Oracle Enterprise Manager Database Control Ports" on page E-3 explains how to modify its port number.	1158	5500-5519	HTTP
Oracle Enterprise Manager Database Control RMI port for Enterprise Manager Database Control. It is configured during installation. "Changing the Oracle Enterprise Manager Database Control Ports" on page E-3 explains how to modify its port number.	5520	5520-5539	TCP
Oracle Enterprise Manager Database Control JMS port for Enterprise Manager Database Control. It is configured during installation. "Changing the Oracle Enterprise Manager Database Control Ports" on page E-3 explains how to modify its port number.	5540	5540-5559	TCP
Enterprise Manager Database Control Agent HTTP port for Enterprise Management Agent. It is configured during installation. "Changing the Oracle Enterprise Management Agent Port" on page E-3 explains how to modify its port number.	3938	1830-1849	HTTP

Table E-1 (Cont.) Ports Used in Oracle Components

Component and Description	Default Port Number	Port Range	Protocol
Oracle XML DB The Oracle XML DB HTTP port is used if Web-based applications need to access an Oracle database from an HTTP listener. It is configured during installation, but you cannot view it afterward. "Changing the Oracle XML DB Ports" on page E-4 explains how to change its port number.	Dynamic	Dynamic	HTTP
Oracle XML DB The Oracle XML DB FTP is used when applications need to access an Oracle database from an FTP listener. It is configured during installation, but you cannot view it afterward. "Changing the Oracle XML DB Ports" on page E-4 explains how to change its port number.	Dynamic	Dynamic	FTP
Oracle RAC (UNIX) The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	UDP
Oracle Clusterware Oracle Clusterware Daemon internode connection. The port number is assigned automatically during installation. You cannot view or modify it afterward.	49896	49896	TCP
Cluster Synchronization Service (CSS) CSS daemon internode connection for the GM layer. The port number is assigned automatically during installation. You cannot view or modify it afterward.	49895	49895	TCP
Oracle Cluster Registry The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	TCP
Oracle Event Manager The port number is assigned automatically during installation. You cannot view or modify it afterward.	49897	49897–49898	TCP
Cluster Manager The port number is assigned automatically during installation. You cannot view or modify it afterward.	Dynamic	Dynamic	TCP

Changing the Oracle Enterprise Management Agent Port

To find the current setting for the Oracle Management Agent port, search for `EMD_URL` in the `$ORACLE_HOME/host_sid/sysman/config/emd.properties` file.

To change the Oracle Management Agent HTTP port, use the `emca -reconfig ports` command:

```
emca -reconfig ports -AGENT_PORT 1831
```

Changing the Oracle Enterprise Manager Database Control Ports

To find the current HTTP, RMI, and JMS port settings, search in the following files:

- HTTP port: Search for `REPOSITORY_URL` in the `$ORACLE_HOME/host_sid/sysman/config/emd.properties` file.

- **RMI port:** Search for the `port` attribute in the `rmi-server` tag in the `$ORACLE_HOME/oc4j/j2ee/OC4J_DBConsole_host_sid/config/rmi.xml` file.
- **JMS port:** Search for the `port` attribute in the `jms-server` tag in the `$ORACLE_HOME/oc4j/j2ee/OC4J_DBConsole_host_sid/config/jms.xml` file.

To change the Oracle Enterprise Manager Database Control ports, use the `emca -reconfig ports` command:

```
$ORACLE_HOME/bin> emca -reconfig ports option setting
```

where *option* can be:

- **DBCONTROL_HTTP_PORT:** Sets the HTTP port, for example:

```
emca -reconfig ports -DBCONTROL_HTTP_PORT 1820
```
- **RMI_PORT:** Sets the RMI port, for example:

```
emca -reconfig ports -RMI_PORT 5520
```
- **JMS_PORT:** Sets the JMS port, for example:

```
emca -reconfig ports -JMS_PORT 5521
```

You can enter multiple `-reconfig port` settings in one line, for example:

```
emca -reconfig ports -DBCONTROL_HTTP_PORT 1820 -AGENT_PORT 1821 -RMI_PORT 5520
```

Changing the Oracle XML DB Ports

To change the Oracle XML DB FTP and HTTP ports, you must run the `catxdbdbca.sql` script, which in a default installation is located in `$ORACLE_HOME/rdbms/admin`.

To change the Oracle XML DB ports:

1. Check that the Oracle listener is running. To do so, in the Services control panel, make sure that the Oracle TNS Listener service (for example, `OracleOraDb11g_home1TNSListener`) is set to **Started**.

If you cannot start the listener, refer to *Oracle Database Net Services Administrator's Guide*.

2. Log into SQL*Plus as SYS or XDB using the SYSDBA role.

For example, to log in to SQL*Plus as SYS using the password `welcome`:

```
$ $ORACLE_HOME/bin/sqlplus
SQL> CONNECT SYS as SYSDBA
Enter password: SYS_password
```

3. Run the `catxdbdbca.sql` script.

For example, to use 2200 for the FTP port and 8200 for the HTTP port, and assuming the Oracle home is in the following location, enter the following command:

```
SQL> $ORACLE_HOME/rdbms/admin/catxdbdbca.sql 2200 8200
```

4. Exit SQL*Plus.

Configuring Oracle Database Globalization Support

This appendix describes the following Globalization Support topics:

- [Installing and Using Oracle Components in Different Languages](#)
- [Running Oracle Universal Installer in Different Languages](#)

Installing and Using Oracle Components in Different Languages

This section describes the following procedures:

- [Configuring Oracle Components to Run in Different Languages](#)
- [Installing Translation Resources](#)

Configuring Oracle Components to Run in Different Languages

You can specify the language and the territory, or locale, in which you want to use Oracle components. The locale setting of a component determines the language of the user interface of the component and the globalization behavior, such as date and number formatting. Depending on the Oracle component, the locale of the component is either inherited from the operating system session that started the component, or is defined by the `NLS_LANG` environment variable.

The operating system locale usually influences Oracle components that are based on Java technology. The `NLS_LANG` environment variable usually influences Oracle components that use Oracle Client libraries such as OCI.

Note: The user interface of an Oracle component will be displayed in a selected language only if the appropriate translation is available and has been installed. Otherwise, the user interface will be displayed in English.

This section describes the following procedures:

- [Determining the Operating System Locale by Using the LANG Environment Variable](#)
- [Configuring Locale and Character Sets by Using the NLS_LANG Environment Variable](#)

Determining the Operating System Locale by Using the LANG Environment Variable

The locale setting of your operating system session determines the language of the user interface and the globalization behavior for components such as Oracle Universal Installer, Oracle Net Configuration Assistant, and Oracle Database Configuration Assistant. It also determines the globalization behavior of Oracle Database sessions created by a user application through Oracle JDBC driver, unless overridden by the application.

The operating system locale on Linux is determined by the value of the `LANG` environment variable. Depending on your desktop environment, such as KDE, GNOME, or telnet, you can select a default session locale on a login screen, in a configuration panel, or in a configuration file.

Note: Refer to the operating system documentation on how to select a locale for the operating system session in your desktop environment.

You can modify the `LANG` variable in the environment of your shell to start an Oracle component in a selected language. For example, to start Oracle Database Configuration Assistant in German, enter one of the following commands:

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

```
$ LANG=de_DE.iso88591 dbca
```

- C shell (csh):

```
% (setenv LANG de_DE.iso88591; dbca)
```

Note: The `LC_ALL` environment variable overrides the value of the `LANG` environment variable. For the commands listed in the following section to work, either ensure that the `LC_ALL` environment variable is not set in the environment, or substitute `LC_ALL` for `LANG`.

To modify the operating system locale for all Oracle components started by the given shell, modify the `LANG` variable using one of the following commands:

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

```
$ LANG=de_DE.iso88591; export LANG  
$ ...
```

- C shell (csh):

```
% setenv LANG de_DE.iso88591  
$ ...
```

The value of the `LANG` environment variable must be a valid operating system locale. To see the list of valid locales, enter the following command:

```
$ locale -a
```

Note: Refer to the operating system documentation for a mapping between values of the `LANG` environment variable and the languages and territories that they represent.

Configuring Locale and Character Sets by Using the NLS_LANG Environment Variable

The NLS_LANG environment variable determines the language of the user interface and the globalization behavior for components such as SQL*Plus, exp, and imp. It sets the language and territory used by the client application and the database user session. It also declares the character set for entering and displaying data by the client application.

The NLS_LANG environment variable uses the following format:

```
NLS_LANG=language_territory.characterset
```

In this format:

- *language* specifies the language used for displaying Oracle messages, sorting, day names, and month names
- *territory* specifies the conventions for default date, monetary, and numeric formats
- *characterset* specifies the encoding used by the client application

In most cases, this is the Oracle character set that corresponds to the character set of the user terminal or the operating system.

The NLS_LANG environment variable is set as a local environment variable for the shell on all UNIX-based platforms. For example, if the operating system locale setting is `en_US.UTF-8`, then the corresponding value of NLS_LANG environment variable is `AMERICAN_AMERICA.AL32UTF8`.

See Also: *Oracle Database Globalization Support Guide* for information about the NLS_LANG parameter and Globalization Support initialization parameters

The following examples illustrate some of the valid values for the NLS_LANG environment variable.

Note: Refer to the operating system documentation on how to determine the operating system locale environment setting.

Operating System Locale	NLS_LANG Values
French (France)	FRENCH_FRANCE.WE8ISO8859P15 FRENCH_FRANCE.WE8ISO8859P1 FRENCH_FRANCE.WE8MSWIN1252 FRENCH_FRANCE.AL32UTF8
Japanese (Japan)	JAPANESE_JAPAN.JA16EUC JAPANESE_JAPAN.JA16SJIS JAPANESE_JAPAN.AL32UTF8

Installing Translation Resources

To view the user interface of Oracle components in different languages, you must install the appropriate language translations along with the component.

Note: Part of Oracle Database Vault user interface text is stored in database tables in the DVSYS schema. By default, only the English language is loaded into these tables. You can use Oracle Database Vault Configuration Assistant to add more languages to Oracle Database Vault. For the necessary steps, refer to Appendix C in *Oracle Database Vault Administrator's Guide*.

To select the translation resources that you want to install:

1. Start Oracle Universal Installer.
2. On the Select Installation Option screen, select the installation option and click **Next**.
3. On the System Class screen, select the type of system class for installing the database, and click **Next**.
4. On the Grid Installation Options screen, select the type of database installation you want to perform, and click **Next**.
5. On the Select Product Languages screen, select the language in which you want to use Oracle components from the Available Languages field.

Note: The Available Languages field lists all languages supported by Oracle globalization libraries. The set of languages for which a translation is actually available is usually smaller and depends on a particular component. The scope of translation for a given component may differ between languages. For example, some translations may include all user interface text, while others may include only error messages and no help files.

6. Use the > arrow to move the selected language to the Selected Languages field, and then click **Next**.

Note: Oracle Universal Installer will ignore languages in the Selected Languages field for which no translation is available.

Running Oracle Universal Installer in Different Languages

Your operating system locale determines the language in which Oracle Universal Installer runs. Oracle Universal Installer may run in one of the following languages:

- Brazilian Portuguese (pt_BR)
- French (fr)
- German (de)
- Italian (it)
- Japanese (ja)
- Korean (ko)
- Simplified Chinese (zh_CN)
- Spanish (es)

- Traditional Chinese (zh_TW)

To run Oracle Universal Installer in one of the available languages, change the locale in which your operating system session is running before you start Oracle Universal Installer with the `./runInstaller` command. If the selected language is not one of them listed earlier, Oracle Universal Installer runs in English.

You must ensure that the selected value for the `LANG` environment variable starts with the appropriate language abbreviation. In the aforementioned list of languages, in which Oracle Universal Installer can run, the required abbreviation appears in parentheses beside the language name. For example, `fr_FR` and `fr_CA` are valid values to run the Oracle Universal Installer in French.

Troubleshooting

This appendix contains information about troubleshooting. It includes information about the following topics:

- [Verify Requirements](#)
- [X Window Display Errors](#)
- [Remote Terminal Installation Error](#)
- [What to Do If an Installation Error Occurs?](#)
- [Reviewing the Log of an Installation Session](#)
- [Troubleshooting Host Name Changes and CSS](#)
- [Troubleshooting Configuration Assistants](#)
- [Troubleshooting Inventory Issues](#)
- [Troubleshooting Screen Display Issues](#)
- [Silent-Mode Response File Error Handling](#)
- [Cleaning Up After a Failed Installation](#)

See Also: Chapter 6, "Troubleshooting Oracle Configuration Manager" in *Oracle Configuration Manager Installation and Administration Guide* for information on some of the errors that may occur while using Oracle Configuration Manager and tips to troubleshoot these errors

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 preinstallation tasks specified in [Chapter 2](#).

Read the Release Notes

Read the release notes for the product before installing it. The release notes are available on the Oracle Database 11g DVD. The latest version of the release notes is also available on the Oracle Technology Network Web site:

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

X Window Display Errors

If you are running Oracle Universal Installer on a remote system and you want to display Oracle Universal 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, refer to 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 fully_qualified_remote_host_name
```

For example:

```
$ xhost somehost.us.example.com
```

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 the X clock and start Oracle Universal Installer again.

Remote Terminal Installation Error

If you run the installation from a remote terminal, or if you use an `su` command to change users you might receive an error similar to the following:

```
Could not execute auto check for display colors using command  
/usr/X11R6/bin/xdpyinfo
```

This can occur if the `DISPLAY` variable is not set, or the user running the installation is not authorized to open an X window. For instance, if you use an `su` command to change from a user that is authorized to open an X window to a user account that is

not authorized to open an X window on the display, such as a lower-privileged user opening windows on the `root` user's console display.

To troubleshoot this issue, run the command `echo $DISPLAY` to ensure that the display variable is set to the correct visual or to the correct host. If the display variable is set correctly then either ensure that you are logged in as the user authorized to open an X window, or run the command `xhost +` to allow any user to open an X window.

What to Do If an Installation Error Occurs?

If you encounter an error during installation:

- Do not exit Oracle Universal Installer.
- If you click **Next** after you enter 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 Oracle Universal Installer is copying or linking files, refer to "[Reviewing the Log of an Installation Session](#)" section on page G-3.
- If you encounter an error while a configuration assistant is running, refer to "[Troubleshooting Configuration Assistants](#)" section on page G-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 G-6.

Reviewing the Log of an Installation Session

During an installation, Oracle Universal 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 the following command to determine the location of the `oraInventory` directory:

```
$ cat /etc/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 Oracle Universal 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 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 Oracle Universal Installer or listed in the log file indicates a relinking problem, refer to the following file for more information:

`$ORACLE_HOME/install/make.log`

Troubleshooting Host Name Changes and CSS

If you change the host name for Automatic Storage Management, then the Oracle CSS daemon will not start. In order to counter this problem, please use the following steps:

- Login as the root user
- Run `localconfig delete` to deconfigure CSS. This will remove any configuration related files on the system that referenced the old host name.
- Run `localconfig add` to reconfigure CSS using the new host name.

For Example:

```
# $ORACLE_HOME/bin/localconfig [add] [delete] [ reset destination_Oracle_home ]
[-silent] [-paramfile Complete_path_of_file_specifying_parameter_values]
```

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 G-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 ["Irrecoverable Errors"](#) on page G-5 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 canceled	-1

Irrecoverable Errors

If you receive an irrecoverable 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 G-6.
2. Correct the cause of the irrecoverable error.
3. Reinstall the Oracle software.

Troubleshooting Inventory Issues

If you face any of the following situations for Oracle home, then run the `opatch lsinventory -detail` command to list the contents of the inventory and see section "Recovering from inventory corruption" in the *Oracle Universal Installer and OPatch User's Guide for Windows and UNIX* for information on fixing the issue.

- Oracle home is cloned without completing the inventory steps.
- There is bad inventory.
- Inventory is not available but it is created when the Oracle Enterprise Manager Agent is installed in a separate Oracle home.

Troubleshooting Screen Display Issues

If you connect to Oracle database with a screen resolution of 640X480 or 800X600, then the Next button in the GUI is not visible as it hides behind the Taskbar. To fix this problem, perform one of the following:

- Hide the Taskbar.
- Move the Oracle Universal Installer screen up.
- Set the screen resolution to 1024X768 or higher.

Silent-Mode Response File Error Handling

To determine whether a silent-mode installation succeeds or fails, refer to the following log file:

```
/oraInventory_location/logs/silentInstalldate_time.log
```

If necessary, refer to 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

For example, a common problem is that while all the product-specific data is filled out correctly, the staging area location may be incorrect. If this is the case, check the `FROM_LOCATION` variable and make sure that it points to the `products.xml` file in the installation media. In the installation media, this `products.xml` is in `response/stage`.

- Oracle Universal Installer encounters an error, such as insufficient disk space

Oracle Universal Installer or configuration assistant validates the response file at run time. If the validation fails, the silent-mode installation or configuration process ends. Oracle Universal 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 Oracle Universal Installer created during the attempted installation and remove the Oracle home directory. Perform the following steps to remove the files:

1. Start Oracle Universal Installer as described in the "[Installing the Oracle Database Software](#)" on page 4-8.
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.

To reinstall, you must drop either one or two database schemas, depending upon the installation type.

Frequently Asked Questions About Installation

Use the following guidelines to decide how to install Oracle Database components:

- [Installing Oracle Database or Oracle Database Client](#)
- [Installing Oracle Database Tools](#)
- [Installing Oracle Database with Oracle Applications](#)
- [Installing Oracle Database Heterogeneous Connectivity Tools \(Gateways\)](#)

Note: Some Oracle Database components may not be available on all platforms. Consult your platform-specific installation guide or release notes.

Installing Oracle Database or Oracle Database Client

The following are frequently asked questions with respect to installing Oracle database:

- [I only need one instance of Oracle Database or I just want to install a test database to get familiar with the product. How do I install Oracle Database for these situations?](#)
- [How can I create an Oracle database that can handle transaction-heavy or data warehousing applications?](#)
- [What's the best way to install multiple Oracle databases?](#)
- [How do I configure client connections to an Oracle database?](#)
- [What is the best way to install Oracle Client if my client nodes have limited disk space?](#)
- [How do I upgrade Oracle Database?](#)
- [The computers at my site have been configured to run as a cluster. How should I install Oracle Database?](#)
- [How do I migrate my non-Oracle databases to Oracle Database?](#)

I only need one instance of Oracle Database or I just want to install a test database to get familiar with the product. How do I install Oracle Database for these situations?

- If you want a quick installation using the default installation settings, then refer to the platform-specific *Oracle Database Quick Installation Guide*.
- If your site has special requirements, then refer to this guide for more information.

How can I create an Oracle database that can handle transaction-heavy or data warehousing applications?

If you want to create a starter database designed for transaction-heavy or data warehousing applications, then refer to this guide for more details. Select the **Advanced Installation** method, and then select the database type you want on the Select Database Configuration screen.

See Also: *Oracle Database Data Warehousing Guide* after installation

Alternatively, you can install Oracle OLAP during the Oracle Database installation. Oracle OLAP provides optimal support for database environments that must meet OLAP requirements. To do so, in the Select Database Edition screen, select **Enterprise Edition**. Click the **Select Options** button, and from the Choose Components screen, select **Oracle OLAP**.

See Also:

- *Oracle OLAP User's Guide*
- *Oracle OLAP DML Reference*
- *Oracle OLAP Java API Reference*

What's the best way to install multiple Oracle databases?

Use this guide to install Oracle Database using either of the following methods:

- **Installing with response files:** This method lets you run Oracle Universal Installer at a command line using a response file that contains settings specific to each computer.
- **Cloning a Database:** Install Oracle Database on one computer using interactive mode. You can also clone databases. Instructions for cloning databases are described in *Oracle Database Administrator's Guide*.

How do I configure client connections to an Oracle database?

1. Install Oracle Database on a server by using this guide for more information.
2. Use platform-specific *Oracle Database Client Installation Guide* to install Oracle Client on each client node, and select the Instant Client installation type.

If you have many client nodes, consider staging the software centrally, mapping the drive, and running Oracle Universal Installer in the silent or response file mode.

If the client nodes only require a default installation into a new Oracle home directory, consider using this guide for more information.

What is the best way to install Oracle Client if my client nodes have limited disk space?

1. Install Oracle Database onto a server by using this guide for more details.

2. Use platform-specific *Oracle Database Client Installation Guide* to install Oracle Client on each client node, and select the Instant Client installation type.

If you have many client nodes, then consider running Oracle Universal Installer in silent or response file mode.

How do I upgrade Oracle Database?

Refer to *Oracle Database Upgrade Guide*.

See Also: *Oracle Database Administrator's Guide* if you want to use software cloning to upgrade Oracle Database

The computers at my site have been configured to run as a cluster. How should I install Oracle Database?

Use any of the following installation scenarios:

- If you want to run a single-instance Oracle Database in a clustered environment, then install Oracle Clusterware either before or after you install Oracle Database.
- If you want a consolidated pool of storage for all databases in a cluster, then install Oracle Clusterware first and use Automatic Storage Management to manage this storage. Afterward, install Oracle Database (which can be either a single instance database or Real Application Clusters).
- If you plan to use Oracle Real Application Clusters, first install Oracle Clusterware, and then install Oracle Real Application Clusters.

Refer to platform-specific *Oracle Grid Infrastructure Installation Guide* and *Oracle Real Application Clusters Installation Guide* for the platform to install Oracle Clusterware or Oracle Real Application Clusters. Oracle Clusterware is available on the Oracle Clusterware installation media. Refer to this guide which explains how to install Automatic Storage Management as well as Oracle Database.

Oracle Clusterware is a key component required by Oracle Real Application Clusters installations. Oracle Clusterware is an integrated cluster management solution that can bind multiple servers together to act as a single system. This is referred to as a cluster. It performs workload management and component restart. For example, when an instance supporting a particular service fails, Oracle Clusterware restarts the service on the next available instance that you have configured for that service. Oracle Clusterware can monitor non-Oracle programs, as long as they are defined within the Oracle Clusterware environment using the High Availability API.

How do I migrate my non-Oracle databases to Oracle Database?

Use Oracle Migration Workbench to migrate your non-Oracle databases and applications to Oracle. Oracle Migration Workbench software and documentation are available at:

<http://www.oracle.com/technology/tech/migration/index.html>

Installing Oracle Database Tools

The following are frequently asked questions with respect to installing Oracle database tools:

- [How do I install Oracle Application Server?](#)
- [How can I administer and monitor my Oracle Database products?](#)
- [How do I manage security for my Oracle Database products?](#)

- How do I use Oracle Database to manage my XML data?
- Does Oracle Database provide OLAP tools so that I can analyze data such as trends and time series in my database?
- Does Oracle Database provide data mining tools that I can use to discover hidden meaning in my data and predict likely outcomes based on my data?
- How do I perform backup and recovery operations for Oracle Database?
- Is Oracle Workflow included with Oracle Database 11g?
- Is there a migration plan for customers that have built solutions using Oracle Workflow?

How do I install Oracle Application Server?

Refer to *Oracle Application Server Installation Guide*. How you install Application Server depends on whether you already have Oracle Database installed:

- If you do not have Oracle Database installed or you do not want Oracle Application Server to use any of your existing Oracle Databases, then Oracle Universal Installer lets you install a separate Oracle Application Server instance. This database is populated with the metadata that Oracle Application Server must run.
- If you want Oracle Application Server to use an existing Oracle Database, then do the following:
 1. From the Oracle Application Server installation media, run Oracle Application Server Repository Creation Assistant to populate your database with the metadata that Application Server needs.
 2. Install the remaining Oracle Application Server components by following the instructions in the *Oracle Application Server Installation Guide*.

How can I administer and monitor my Oracle Database products?

To perform regular administrative functions such as creating, configuring, or deleting databases, or managing database templates, use one of the following methods:

To manage only the single database and listener that you are installing:

1. Use this guide to install Oracle Database.
2. From Oracle Database, use Database Configuration Assistant to manage your databases.

You can also administer and monitor the database with Oracle Enterprise Manager Grid Control, which is installed by default with Oracle Database. Oracle Enterprise Manager Grid Control requires an agent which is not installed by default.

Oracle Enterprise Manager Grid Control includes the Oracle Management Agent, Oracle Management Service, and Oracle Management Repository, as well as Grid Control, a browser-based central console through which administrators can perform all monitoring, administration, and configuration tasks for the enterprise.

See Also: *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* available on the Enterprise Manager Grid Control installation media

Documentation available on the Oracle Technology Network Web site at:

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

To perform advanced administration tasks, such as monitoring Oracle Database and managing multiple hosts, application servers, and databases including the one that you are installing, install Oracle Enterprise Manager as follows:

1. Use this guide to install Oracle Database.

If you plan to use Oracle Real Application Clusters, then install Oracle Database by using platform-specific *Oracle Grid Infrastructure Installation Guide* and *Oracle Real Application Clusters Installation Guide*.

2. Use *Oracle Enterprise Manager Grid Control Installation and Basic Configuration* to install and configure Oracle Enterprise Manager. For postconfiguration tasks, use *Oracle Enterprise Manager Advanced Configuration*. Refer to documentation available on the Enterprise Manager Grid Control installation media, or on the Oracle Technology Network Web site at:

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

How do I manage security for my Oracle Database products?

Oracle provides a wide range of security solutions for your enterprise environment, including centralized administration and security features integrated with Oracle Internet Directory. The set of Oracle security services called Oracle Platform Security integrates the security features built into Oracle Database, Oracle Application Server, and the Oracle Identity Management infrastructure. Combined, these features enable the development and deployment of secure e-business applications.

Oracle Identity Management includes Oracle Internet Directory, a centralized repository that simplifies administration of users and applications in the Oracle environment by means of the following components:

- Oracle Internet Directory client tools, including LDAP command-line tools, the Oracle Internet Directory SDK, and Oracle Directory Manager.
- Oracle Internet Directory server components, including the directory server, the directory replication server, the directory integration server, and various tools for starting and stopping them.

Oracle Database includes the Oracle Internet Directory client tools, but not the Oracle Internet Directory server components. To install the Oracle Internet Directory server components, run Oracle Universal Installer from an Oracle 10g Application Server installation.

See Also:

- *Oracle Application Server Installation Guide* (to install Oracle Identity Management)
- *Oracle Database Security Guide*
- *Oracle Database Advanced Security Administrator's Guide*
- *Oracle Database Enterprise User Security Administrator's Guide*
- *Oracle Label Security Administrator's Guide*
- *Oracle Application Server Security Guide*
- Oracle Technology Network topics on database security
<http://www.oracle.com/technology/deploy/security/index.html>

How do I use Oracle Database to manage my XML data?

Use Oracle XML DB, which is installed as part of Oracle Database. Oracle XML DB enables you to efficiently store, generate, retrieve, query, and manage XML data on your site. Oracle XML DB provides all the advantages of a relational database, for example, allowing you to control the referential integrity of XML data with constraints and triggers. It works well with large amounts of XML data by storing it in a parsed, relational form, which improves access performance.

Oracle XML DB supports XML Type, which is a native data type for XML data, for which you can choose various storage options depending on your needs. In addition, Oracle XML DB supports XML Schema processing, structured and unstructured storage, a content repository that you can access by using common protocols (FTP, HTTP(S), and WebDAV), and SQL/XML, which is a standard for SQL with XML. For Oracle Database 11g Release 1 (11.1), Oracle XML DB introduced support for the XQuery language for querying, transforming, and constructing XML; the ability for users to define their own metadata for schema-based XML; a set of new SQL functions for DML operations on XML data; and more.

You can use Oracle XML DB in conjunction with Oracle XML Developer's Kit (XDK) to build applications that run on either Oracle Database or Oracle Application Server.

See Also:

- *Oracle XML DB Developer's Guide*
- *Oracle XML Developer's Kit Programmer's Guide*

Does Oracle Database provide OLAP tools so that I can analyze data such as trends and time series in my database?

Yes, install Oracle OLAP, which is provided in the Oracle Database installation. Oracle OLAP provides optimal support for database environments that must meet OLAP requirements.

Use either of the following methods in *Oracle Database Installation Guide* to install Oracle OLAP:

- When you run Oracle Universal Installer, in the Select Database Edition screen, select **Enterprise Edition**. Click the **Select Options** button, and from the Choose Components screen, select **Oracle OLAP**.

See Also:

- *Oracle OLAP User's Guide*
 - *Oracle OLAP DML Reference*
 - *Oracle OLAP Java API Reference*
- Select the **Enterprise Edition** installation type, and then on the Select Database Configuration screen, select the **Data Warehouse** configuration.

See Also: *Oracle Database Data Warehousing Guide* after installation

Does Oracle Database provide data mining tools that I can use to discover hidden meaning in my data and predict likely outcomes based on my data?

Yes, you must have an Enterprise Edition licence for the database installation. Install Oracle Data Mining, which is provided in the Oracle Database installation. With the Oracle Data Mining option, you can create and execute predictive and descriptive data mining models that use a variety of algorithms.

Use the following method in this guide to install Oracle Data Mining:

1. When you run Oracle Universal Installer, select the **Enterprise Edition** installation type.
2. In the Select Database Configuration screen, select the **General Purpose/Transaction Processing** configuration.

See Also: The following manuals after you have installed Oracle Data Mining:

- *Oracle Data Mining Concepts*
- *Oracle Data Mining Administrator's Guide*
- *Oracle Data Mining Application Developer's Guide*
- *Oracle Data Mining Java API Reference*
- *Oracle Database PL/SQL Packages and Types Reference* (search for Data Mining)

How do I perform backup and recovery operations for Oracle Database?

Use Oracle Database Recovery Manager (RMAN), which is a backup and recovery tool integrated into Oracle Database. This tool satisfies the pressing demands of high-performance, manageable backup, and recovery. Recovery Manager is native to the database server, automatically tracks database structure changes, and optimizes operations accordingly. In addition, Recovery Manager is integrated with leading tape media management products, so that Oracle database backups can be integrated with your existing networked data protection infrastructure.

See Also:

- *Oracle Database Backup and Recovery User's Guide*
- *Oracle Database Backup and Recovery Reference*

Is Oracle Workflow included with Oracle Database 11g?

Starting with Oracle Database 11g, Oracle Workflow is no longer released with the database. Oracle Workflow will be available with the Oracle E-Business Suite releases.

See Also: Oracle Workflow statement of direction

http://www.oracle.com/technology/products/ias/workflow/workflow_sod.html

Is there a migration plan for customers that have built solutions using Oracle Workflow?

Starting January 2006, customers are encouraged to re-create and implement workflows using Oracle BPEL Process Manager. Oracle is in the process of creating a technical migration guide that will provide detailed recommendations for migrating Oracle Workflow processes to Oracle BPEL Process Manager.

See Also: Oracle Workflow statement of direction

http://www.oracle.com/technology/products/ias/workflow/workflow_sod.html

Installing Oracle Database with Oracle Applications

The following are frequently asked questions with respect to installing Oracle database with Oracle applications:

- [How do I install my Oracle applications with Oracle Database?](#)
- [How can I create Web applications that communicate with Oracle Database?](#)
- [Which Web server can my Oracle applications use?](#)
- [How can I migrate my non-Oracle applications to Oracle?](#)

How do I install my Oracle applications with Oracle Database?

In most cases, install Oracle Database itself, then install the Oracle application. The Oracle Universal Installer for that application prompts you for the connection information. Check the application documentation requirements.

If you must implement your applications with Oracle Real Applications Clusters databases, refer to *Oracle Real Application Clusters Installation Guide* and *Oracle Grid Infrastructure Installation Guide* for more information.

How can I create Web applications that communicate with Oracle Database?

Install Oracle Application Express and a Web server.

Use this guide to install Oracle Database. Oracle Application Express is automatically installed, when you install Oracle database.

See Also: *Oracle Application Express Installation Guide*

Which Web server can my Oracle applications use?

Install Oracle HTTP Server:

Use this guide to install Oracle Database.

How can I migrate my non-Oracle applications to Oracle?

Use Oracle Migration Workbench to migrate your non-Oracle applications to Oracle. Oracle Migration Workbench software and documentation are available at:

<http://www.oracle.com/technology/tech/migration/index.html>

Installing Oracle Database Heterogeneous Connectivity Tools (Gateways)

The following section discusses the Gateway products:

[How can my Oracle applications access data in a non-Oracle database system?](#)

How can my Oracle applications access data in a non-Oracle database system?

You can use Oracle Database Gateway as the connectivity tool to enable Oracle applications to access data in non-Oracle databases. The following are the functions of Oracle Database Gateway:

- Integrates a non-Oracle database into your Oracle Database environment.
- Enables Oracle PL/SQL applications to integrate with APPC-enabled transactions, or access messages in IBM Websphere MQ.

You can install the Gateway product on a computer independent of the Oracle application, Oracle database, and non-Oracle database.

For example, suppose you have the following scenario:

- Oracle Database is installed on an UNIX computer.
- The Oracle application is installed on a Microsoft Windows computer and accesses data from the Oracle database on the UNIX computer.
- The Oracle application must join data in a DB2 database on Solaris Operating System and an Oracle Database on UNIX.

You have the option of installing the Database Gateway for DRDA on the Solaris computer where DB2 is running, on UNIX where Oracle is running, or on a third computer.

[Table H-1](#) lists the non-Oracle database systems that you can access from Oracle applications, and the Gateways products that are available for those systems.

Table H-1 Oracle Gateway Products

Non-Oracle Database	Oracle Gateway Products and Documentation
IBM DB2 Universal Database (UDB)	Oracle Database Gateway for DRDA. <i>Use Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64 and Oracle Database Gateway for DRDA User's Guide.</i>
IBM DB2 z/OS	Oracle Database Gateway for DRDA. <i>Use Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64 and Oracle Database Gateway for DRDA User's Guide.</i>
IBM DB2/400	Oracle Database Gateway for DRDA. <i>Use Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64 and Oracle Database Gateway for DRDA User's Guide.</i>
WebSphere MQ	Oracle Database Gateway for WebSphere MQ. <i>Oracle Database Gateway for WebSphere MQ Installation and User's Guide.</i>
CICS/TS IMSTM	Oracle Database Gateway for APPC. <i>Use Oracle Database Gateway for APPC Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64.</i> <i>Use Oracle Database Gateway for APPC User's Guide</i>

Table H-1 (Cont.) Oracle Gateway Products

Non-Oracle Database	Oracle Gateway Products and Documentation
SQL Server	<p>Oracle Database Gateway for SQL Server.</p> <p>Use <i>Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i> and <i>Oracle Database Gateway for SQL Server User's Guide</i>.</p>
Sybase Adaptive Server	<p>Oracle Database Gateway for Sybase.</p> <p>Use <i>Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i> and <i>Oracle Database Gateway for Sybase User's Guide</i>.</p>
Teradata	<p>Oracle Database Gateway for Teradata.</p> <p>Use <i>Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i> and <i>Oracle Database Gateway for Teradata User's Guide</i>.</p>
Informix Server	<p>Oracle Database Gateway for Informix.</p> <p>Use <i>Oracle Database Gateway Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i> and <i>Oracle Database Gateway for Informix User's Guide</i>.</p>
IMS	<p>Oracle Database Gateway for IMS.</p> <p>Use <i>Oracle Database Gateway for IMS, VSAM, and Adabas Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i>, <i>Oracle Database Gateway for IMS User's Guide</i> and <i>Oracle Connect for IMS, VSAM, and Adabas Gateways Installation and Configuration Guide for IBM z/OS</i></p>
VSAM	<p>Oracle Database Gateway for VSAM.</p> <p>Use <i>Oracle Database Gateway for IMS, VSAM, and Adabas Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i>, <i>Oracle Database Gateway for VSAM User's Guide</i> and <i>Oracle Connect for IMS, VSAM, and Adabas Gateways Installation and Configuration Guide for IBM z/OS</i>.</p>
Adabas	<p>Oracle Database Gateway for Adabas.</p> <p>Use <i>Oracle Database Gateway for IMS, VSAM, and Adabas Installation and Configuration Guide for AIX 5L Based Systems (64-Bit), HP-UX Itanium, Solaris Operating System (SPARC 64-Bit), Linux x86, and Linux x86-64</i>, <i>Oracle Database Gateway for Adabas User's Guide</i> and <i>Oracle Connect for IMS, VSAM, and Adabas Gateways Installation and Configuration Guide for IBM z/OS</i>.</p>

Country Codes

This appendix contains a list of valid country codes that can be used while installing Oracle Configuration Manager.

Valid Country Codes

[Table I-1](#) contains a list of countries and their short names (codes).

Table I-1 Country Codes

Country	Short Name (Code)
African Other	AA
Andorra	AD
United Arab Emirates	AE
Afghanistan	AS
Antigua and Barbuda	AM
Anguilla	AI
Albania	AL
Armenia	AM
Netherlands Antilles	AN
Angola	AO
Antarctica	AQ
Argentina	AR
American Samoa	AS
Austria	AT
Australia	AU
Aruba	AW
Azerbaijan	AZ
Bosnia-Herzegovina	BA
Barbados	BB
Bangladesh	BD
Belgium	BE
Burkina Faso	BF

Table I-1 (Cont.) Country Codes

Country	Short Name (Code)
Bulgaria	BG
Bahrain	BH
Burundi	BI
Benin	BJ
Bermuda	BM
Brunei Darussalam	BN
Bolivia	BO
Brazil	BR
Bahamas	BS
Bhutan	BT
Bouvet Island	BV
Botswana	BW
Belarus	BY
Belize	BZ
Canada	CA
Cocos (Keeling) Islands	CC
Central African Republic	CF
Congo	CG
Switzerland	CH
Cote D'Ivoire	CI
Cook Islands	CK
Chile	CL
Cameroon	CM
China	CN
Columbia	CO
Costa Rica	CR
Cuba	CU
Cape Verde	CV
Christmas Island	CX
Cyprus	CY
Czech Republic	CZ
Germany	DE
Djibouti	DJ
Denmark and Iceland	DK
Dominica	DM
Dominican Republic	DO
Algeria	DZ

Table I-1 (Cont.) Country Codes

Country	Short Name (Code)
Ecuador	EC
Estonia	EE
Egypt	EG
Western Sahara	EH
Eritrea	ER
Spain	ES
Ethiopia	ET
Finland	FI
Fiji	FJ
Falkland Islands (Malvinas)	FK
Micronesia (Federated States Of)	FM
Faroe Islands	FO
France	FR
France - Overseas Territories	FX
Gabon	GA
United Kingdom	GB
Grenada	GD
Georgia	GE
French Guiana	GF
Ghana	GH
Gibraltar	GI
Greenland	GL
Gambia	GM
Guinea	GN
Guadeloupe	GP
Equatorial Guinea	GQ
Greece	GR
South Georgia and South Sandwich Island	GS
Guatemala	GT
Guam	GU
Guinea - Bissau	GW
Guyana	GY
Hong Kong	HK
Heard Island and McDonald Islands	HM
Honduras	HN
Croatia	HR
Haiti	HT

Table I-1 (Cont.) Country Codes

Country	Short Name (Code)
Hungary	HU
Indonesia	ID
Ireland	IE
Israel	IL
India	IN
British Indian Ocean Territory	IO
Iraq	IQ
Iran (Islamic Republic of)	IR
Iceland	IS
Italy	IT
Jamaica	JM
Jordan	JO
Japan	JP
Kenya	KE
Kyrgyzstan	KG
Cambodia	KH
Kiribati	KI
Comoros	KM
Saint Kitts and Nevis	KN
Democratic People's Republic of Korea	KP
Republic of Korea	KR
Kuwait	KW
Cayman Islands	KY
Kazakhstan	KZ
Lao People's Democratic Republic	LA
Lebanon	LB
Saint Lucia	LC
Liechtenstein	LI
Sri Lanka	LK
Liberia	LR
Lesotho	LS
Lithuania	LT
Luxembourg	LU
Latvia	LV
Libyan Arab Jamahiriya	LY
Morocco	MA
Monaco	MC

Table I-1 (Cont.) Country Codes

Country	Short Name (Code)
Republic of Moldova	MD
Madagascar	MG
Marshall Islands	MH
Macedonia	MK
Mali	ML
Myanmar	MM
Mongolia	MM
Macau	MO
Northern Mariana Islands	MP
Martinique	MQ
Mauritania	MR
Montserrat	MS
Malta	MT
Mauritius	MU
Malawi	MW
Mexico	MX
Malyasia	MY
Mozambique	MZ
Namibia	NA
New Caledonia	NC
Niger	NE
Norfolk Island	NF
Nigeria	NG
Nicaragua	NI
Netherlands	NL
Norway	NO
Nepal	NP
Narau	NR
Niue	NU
New Zealand	NZ
Oman	OM
Panama	PA
Peru	PE
French Polynesia	PF
Papua New Guinea	PG
Philippines	PH
Pakistan	PK

Table I-1 (Cont.) Country Codes

Country	Short Name (Code)
Poland	PL
Saint Pierre and Miquelon	PM
Pitcairn	PN
Puerto Rico	PR
Portugal	PT
Palau	PW
Paraguay	PY
Qatar	QA
Reunion	RE
Romania	RO
CIS-Comm. of Indep. States	RU
Rwanda	RW
Saudi Arabia	SA
Solomon Islands	SB
Seychelles	SC
Sudan	SD
Sweden	SE
Singapore	SG
Saint Helena	SH
Slovenia	SI
Svalbard and Jan Mayen Islands	SJ
Slovakia	SK
Sierra Leone	SL
San Marino	SM
Senegal	SN
Somalia	SO
Suriname	SR
Sao Tome and Principe	ST
El Salvador	SV
South Asia Growth Economies	SX
Syrian Arab Republic	SY
Swaziland	SZ
Turks and Caicos Islands	TC
Chad	TD
French Southern Territories	TF
Togo	TG
Thailand	TH

Table I-1 (Cont.) Country Codes

Country	Short Name (Code)
Tajikistan	TJ
Tokelau	TK
Turkmenistan	TM
Tunisia	TN
Tonga	TO
East Timor	TP
Turkey	TR
Trinidad and Tobago	TT
Tuvalu	TV
Taiwan - Republic of China	TW
United Republic of Tanzania	TZ
Ukraine	UA
Uganda	UG
United States Minor Outlying Islands	UM
United States	US
Uruguay	UY
Uzbekistan	UZ
Vatican City State (Holy See)	VA
Saint Vincent and the Grenadines	VC
Venezuela	VE
Virgin Islands (British)	VI
Vietnam	VN
Vanuatu	VU
Wallis and Futuna Islands	WF
Samoa	WS
Yemen	YE
Mayotte	YT
Serbia and Montenegro	YU
South Africa	ZA
Zambia	ZM
Zaire	ZR
Zimbabwe	ZW

Glossary

Automatic Storage Management disk group

A set of disk devices that **Automatic Storage Management** manages as a single unit. Each disk device can be an individual physical disk, a multiple disk device such as a RAID storage array or logical volume, or even a partition on a physical disk. You can create the Automatic Storage Management disk group when you create the **Automatic Storage Management instance**, or with Oracle Database Configuration Assistant.

Automatic Storage Management instance

The Oracle instance that manages **Automatic Storage Management disk groups**. It is created automatically when you install and configure **Automatic Storage Management**. See also Oracle system identifier (**SID**).

Automatic Storage Management

Enables creation of a single disk group from a collection of individual disk devices. It balances I/O to the disk group across all of the devices in the disk group. It also implements striping and mirroring to improve I/O performance and data reliability.

automatic undo management mode

A mode of Oracle Database in which undo data is stored in a dedicated **undo tablespace**. Unlike in **manual undo management mode**, the only undo management that you must perform is the creation of the undo tablespace. All other undo management is performed automatically.

connect descriptor

A specially formatted description of the destination for a network connection. A connect descriptor contains destination service and network route information.

The destination service is indicated by using its service name for the Oracle Database or its Oracle system identifier (**SID**) for Oracle release 11.2 databases. The network route provides, at a minimum, the location of the **listener** through use of a network address.

connect identifier

A name, net service name, or service name that resolves to a connect descriptor. Users initiate a connect request by passing a user name and password along with a connect identifier in a connect string for the service to which they want to connect, for example:

```
SQL> CONNECT user_name@connect_identifier
Enter password: password
```

control files

Files that record the physical structure of a database and contain the database name, the names and locations of associated databases and online **undo tablespace**, the time stamp of the database creation, the current log sequence number, and checkpoint information.

default domain

The network domain within which most client requests take place. It can be the domain where the client resides, or a domain from which the client often requests network services. The default domain is also the client configuration parameter that determines what domain to append to unqualified network name requests. A name request is unqualified if it does not have a "." character within it.

directory naming

A **naming method** that specifies a directory server to resolve a net service name into a connect descriptor. The net service name is stored centrally in a directory server.

directory server

A Lightweight Directory Access Protocol (LDAP)-compliant directory server. A directory can provide centralized storage and retrieval of database network components, user and corporate policies preferences, user authentication, and security information, replacing client-side and server-side localized files.

external procedures

Procedure or function written in the C programming language and stored in a shared library. An Oracle server can call external procedures or functions using PL/SQL routines. For Oracle Database to connect to external procedures, the server must be configured with a net service name and the **listener** must be configured with protocol address and service information.

global database name

The full database name that uniquely distinguishes it from any other database in your network domain.

For example:

```
sales.us.example.com
```

where `sales` is the name you want to call your database and `us.example.com` is the network domain in which the database is located.

initialization parameter file

An ASCII text file that contains information needed to initialize a database and **instance**.

instance

Process associated with a running Oracle Database instance. When a database is started on a database server (regardless of the type of computer), Oracle Database allocates a memory area called the **System Global Area** and starts one or more Oracle Database processes. This combination of the System Global Area and Oracle Database processes is called an instance. The memory and processes of an instance manage the associated database's data efficiently and serve the users of the database.

installation type

A predefined component set that automatically selects which components to install. See "Oracle Database Editions" on page 1-8 for a list of installation types available with each top-level component.

Interprocess Communication (IPC)

A protocol that client applications use that resides on the same node as the **listener** to communicate with the database. IPC can provide a faster local connection than TCP/IP.

listener

A process that resides on the server and whose responsibility is to listen for incoming client connection requests and manage the traffic to the server.

When a client requests a network session with a database server, a listener receives the actual request. If the client information matches the listener information, then the listener grants a connection to the database server.

listener.ora file

A configuration file for the listener that identifies the:

- Listener name
- Protocol addresses on which it is accepting connection requests
- Services for which it is listening

The `listener.ora` file resides in the `$ORACLE_HOME/network/admin` directory.

An Oracle Database 11g Release 2 (11.2) does not require identification of the database service because of service registration. However, static service configuration is required for an Oracle Database 11g Release 2 (11.2) if you plan to use Oracle Enterprise Manager.

local naming

A **naming method** that resolves a net service name into a connect descriptor. This name is configured and stored in the **tnsnames.ora file** on each individual client.

manual undo management mode

A mode of the database in which undo blocks are stored in user-managed rollback segments.

naming method

A resolution method used by a client application to resolve a connect identifier to a network address when attempting to connect to a database service. Oracle Net Services supports the following naming methods:

- Local naming
- Directory naming
- Host naming
- External naming

net service name

A simple name for a service that resolves to a connect descriptor. Users initiate a connect request by passing a user name and password along with a net service name in a connect string for the service to which they want to connect:

```
SQL> CONNECT user_name@net_service_name
Enter password: password
```

Depending on your needs, net service names can be stored in a variety of places, including:

- Local configuration file, `tnsnames.ora`, on each client
- Directory server
- External naming service, such as Network Information Service (NIS) or Cell Directory Service (CDS)

OPSS\$

Acronym for operating system specific. The initialization file parameter `OS_AUTHENT_PREFIX` enables users to specify a prefix that Oracle uses to authenticate users attempting to connect to the database. Oracle concatenates the value of this parameter to the beginning of the user's operating system account name and password. When a connection request is attempted, Oracle compares the prefixed user name with Oracle user names in the database.

The default value of this parameter is " " (a null string), thereby eliminating the addition of any prefix to operating system account names. In earlier releases, `OPSS$` was the default setting.

ORACLE_BASE

`ORACLE_BASE` is the root of the Oracle Database directory tree. The Oracle Base directory is the top level directory that you can use to install the various Oracle software products. You can use the same Oracle base directory for more than one installation. For example, `/u01/app/oracle` is an Oracle base directory created by the Oracle user.

ORACLE_HOME

Corresponds to the environment in which Oracle Database products run. If you install an OFA-compliant database, using Oracle Universal Installer defaults, Oracle home (known as `$ORACLE_HOME` in this guide) is located beneath `$ORACLE_BASE`. The default Oracle home is `db_n` where *n* is the Oracle home number. It contains subdirectories for Oracle Database software executables and network files. See also [Oracle home](#).

Oracle home

The directory path to install Oracle components (for example, `/u01/app/oracle/product/11.2.0/db_n`). You are prompted to enter an Oracle home in the Path field of the Specify File Locations window. See also [ORACLE_HOME](#), [Oracle home name](#).

Oracle home name

The name of the current Oracle home, for example, `dbhome_1`. Each Oracle home has a home name that distinguishes it from all other Oracle homes on your computer. During installation, you are prompted to enter an Oracle home name in the Name field on the Specify File Locations window.

Oracle schema

A set of rules that determine what can be stored in an LDAP-compliant directory server. Oracle has its own schema that is applied to many types of Oracle entries,

including Oracle Net Services entries. The Oracle schema for Oracle Net Services entries includes the attributes the entries may contain.

Oracle Documentation Library

The media in your kit that includes the Oracle Database documentation. The Oracle Documentation Library is separate from the installation media.

The Oracle Documentation Library does not include this installation guide or *Oracle Database Release Notes for Linux*. These documents are included on the media labeled Oracle Database 11g Release 2 (11.2) and are available on Oracle Technology Network (OTN).

Oracle Net foundation layer

A networking communication layer that establishes and maintains the connection between the client application and server, as well as exchanging messages between them.

protocol address

An address that identifies the network address of a network object.

When a connection is made, the client and the receiver of the request, such as the [listener](#), or Oracle Connection Manager, are configured with identical protocol addresses. The client uses this address to send the connection request to a particular network object location, and the recipient "listens" for requests on this address. It is important to install the same protocols for the client and the connection recipient, and to configure the same addresses.

raw partitions

Portions of a physical disk that are accessed at the lowest possible disk (block) level.

redo log files

Files that contain a record of all changes made to data in the database buffer cache. If an instance failure occurs, then an administrator can use the redo log files to recover the modified data that was in memory.

repository

A set of tables located in any Oracle database accessible to the Oracle Management Server. Oracle Management Server uses a repository to store all system data and application data, information about the state of managed nodes distributed throughout the environment, as well as information about the separately licensable management packs.

service registration

A feature by which the PMON process (an instance background process) automatically registers information with a [listener](#). Because this information is registered with the listener, the [listener.ora file](#) does not need to be configured with this static information.

Service registration provides the listener with the following information:

- Service name(s) for each running instance of the database
- Instance name(s) of the database
- Service handlers (dispatchers and dedicated servers) available for each instance

This enables the listener to direct a client's request appropriately.

- Dispatcher, instance, and node load information

This enables the listener to determine which dispatcher can best handle a client connection's request. If all dispatchers are blocked, the listener can spawn a dedicated server for the connection.

This information enables the listener to determine how best to service a client connection request.

SID

The Oracle system identifier that distinguishes the database from all other databases on your computer. The SID automatically defaults to the database name portion of the global database name (`sales` in the example `sales.us.example.com`) until you reach eight characters or enter a period. You can accept or change the default value.

The SID can also refer to an Automatic Storage Management instance SID, available when you install [Automatic Storage Management](#).

sqlnet.ora file

A configuration file for the client or server that specifies the:

- Client domain to append to unqualified service names or net service names
- Order of naming methods for the client to use when resolving a name
- Logging and tracing features to use
- Route of connections
- External naming parameters
- Oracle Advanced Security parameters

The `sqlnet.ora` file resides in `$ORACLE_HOME/network/admin`.

Secure Sockets Layer (SSL)

An industry standard protocol designed by Netscape Communications Corporation for securing network connections. SSL provides authentication, encryption, and data integrity using public key infrastructure (PKI).

SSL

See [Secure Sockets Layer \(SSL\)](#).

System Global Area

A group of shared memory structures that contain data and control information for an Oracle Database [instance](#).

system identifier

See [SID](#).

tablespace

A logical storage unit within a database. Tablespaces are divided into logical units of storage called segments, which are further divided into extents.

tnsnames.ora file

A configuration file that contains net service names mapped to connect descriptors. This file is used for the local naming method. The `tnsnames.ora` file resides in `$ORACLE_BASE/network/admin`.

undo tablespace

A tablespace that contains one or more undo segments. The creation of any other types of segment (for example, tables, indexes) in undo tablespaces is not allowed.

In the automatic mode, each Oracle instance is assigned one and only one undo tablespace. Each undo tablespace is composed of a set of undo files. Undo blocks are grouped in extents. At any point in time, an extent is either allocated to (and used by) a transaction table, or is free.

Blocks in undo tablespaces are grouped into the following categories:

- File control blocks, bitmap blocks, and so forth used for space management
- Undo segments containing transaction table blocks, undo blocks, and extent-map blocks used for transaction management
- Free blocks that are unallocated to file control or undo segments

unqualified name

A net service name that does not contain a network domain.

Index

Symbols

/, G-3

A

accounts

- reviewing, 6-5
- unauthenticated access to, 6-9

ACFS, 1-11, 3-3

- requirements, 3-3

ADVM, 3-3

- requirements, 3-3

aio-max-nr file, 2-25

aliases, multiple on computers, 2-17

ANONYMOUS user

- unauthenticated account access with, 6-9

APPC-enabled databases, H-9

applications, migrating non-Oracle applications to Oracle, H-8

asm groups

- creating, 2-22

ASM *See* Automatic Storage Management

ASMCA, 6-3

asmcmd utility, 3-19

asmdba

- creating the asmdba group, 2-22

asmdba groups

- creating, 2-22

Automatic Storage Management

- asmcmd utility, 3-19
- characteristics of failure groups, 3-8
- checking disk availability, 3-12
- configuring disks, 3-6 to ??, 3-10
- configuring disks for Automatic Storage Management, 3-10
- considerations before installing, 3-5
- DAS disks, 3-10
- disk devices, 1-12
- disk groups, 3-7
- disks, supported, 3-10
- displaying attached disks, 3-12
- failure groups
 - examples, 3-8
 - identifying, 3-8
- identifying available disks, 3-12

identifying disks, 3-12

installation, testing, 3-19

managing, 6-3

mirroring, 3-7

multiple databases, 2-20

Optimal Flexible Architecture file naming conventions, D-5

partition creation, 3-10

password file, 3-5

recommendations for disk groups, 3-7

redundancy levels, 3-7

response files, A-3

SAN disks, 3-10

space required for preconfigured database, 3-7

SPFILE server parameter file, 3-5

starting and stopping, 6-3

Automatic Storage Management (ASM)

ASM disk group templates, 1-12

templates, 1-12

Automatic Storage Management Configuration

Assistant, 6-3

Automatic Storage Management disk groups

about, 1-12

managing, 6-3

Automatic Storage Management failure groups

about, 1-12

Automatic Storage Management instance

about, 1-12

B

backups of database

Oracle Database Recovery Manager, H-7

base directory

See Oracle base directory

Basic installation type

response file installations, A-6

C

certification, hardware and software, 1-5

checking distribution of the operating system, 2-7

checking version of the operating system, 2-7

chmod command, 2-32, 2-34

chown command, 2-32, 2-34

CLASSPATH environment variable, 5-9

- client static library, generating, 5-4
- cloning
 - Configuring Oracle Configuration Manager in a Cloned Oracle Home, B-3
 - Oracle home, B-1
- Cluster Manager
 - ports, ranges and protocol, E-3
- Cluster Ready Services (CRS). *See* Oracle Clusterware
- Cluster Synchronization Services (CSS)
 - Automatic Storage Management, 1-12
 - ports, ranges and protocol, E-3
- clusters
 - installation guidelines, 4-2
 - See also* Oracle Clusterware, Oracle Real Application Clusters
- Clusterware
 - installed before Oracle Database, 4-2
- Clusterware. *See* Oracle Clusterware
- commands
 - fdisk, 2-35
 - useradd, 2-23
 - usermod, 2-23
- computers with multiple aliases, 2-17
- computers, non-networked, 2-17
- configuration assistants
 - failure, G-4
 - troubleshooting, G-4
- configuring
 - accounts of Oracle users, 5-3
 - kernel parameters, 2-25
- configuring disks for Automatic Storage Management, 3-6 to ??, 3-10
- Configuring Oracle Configuration Manager in a Cloned Oracle Home, B-3
- Connection Manager
 - ports, ranges and protocol, E-2
- control files
 - locating, 6-13
 - naming, D-4
 - reviewing, 6-11
 - using Oracle Enterprise Manager Database Control with, 6-13
- CONTROL_FILES initialization parameter, 6-13
- country codes, I-1
- CSD
 - download location for WebSphere MQ, 2-15
 - requirements
 - on Linux, 2-15
- csh.login file, 2-24
- csh.login.local file, 2-24
- custom database
 - failure groups for Automatic Storage Management, 3-8
 - requirements when using Automatic Storage Management, 3-7
- custom install option, 4-11

D

- DAS (direct attached storage) disks, 3-10
- data files
 - creating separate directories for, 2-34
 - defined, 6-11
 - managing with Automatic Storage Management, 1-11
 - minimum disk space for, 2-33
 - naming, D-4
 - options for placing on file system, 2-33
 - recommendations for file system, 2-33
 - reviewing, 6-11
 - setting permissions on data file directories, 2-34
 - setting up, 6-11
- data loss
 - minimizing with Automatic Storage Management, 3-8
- data mining tools
 - Oracle Data Mining, H-7
- data warehousing tool
 - Oracle OLAP, H-6
- Database Configuration Assistant
 - running in silent mode, A-8
 - troubleshooting, G-4
- databases
 - Automatic Storage Management
 - requirements, 3-7
 - files, 6-11
 - identifying, 6-10
 - initialization parameter file, 6-11
 - naming, 4-12
 - non-Oracle
 - APPC-enabled, H-9
 - non-Oracle, listed, H-9
 - OLAP support (Oracle OLAP), H-6
 - Optimal Flexible Architecture file naming conventions, D-4
 - recovery with Oracle Backup and Recovery, H-7
 - redo log files, 6-12
 - security management, H-5
 - tablespaces, 6-11
- Daylight Savings Time, 1-17
- DB_DOMAIN initialization parameter, 6-10
- DB_NAME initialization parameter, 6-10
- DB2 database, H-9
- DB2 z/OS database, H-9
- DB2/400 database, H-9
- dba group
 - creating, 2-21
 - description, 2-19
 - SYSDBA privilege, 2-19
- dba groups
 - creating, 2-22
- dbca.rsp file, A-5
- default data files, 6-12
- default file mode creation mask
 - setting, 2-38
- default tablespaces, 6-12
- Deinstallation Tool, 7-1
- description
 - database restart, 3-1
 - Oracle Restart, 3-1

- device names
 - IDE disks, 3-12
 - RAID, 3-12
 - SCSI disks, 3-12
- DHCP computers, installing on, 2-16
- directory
 - creating separate data file directories, 2-34
 - database file directory, 2-33
 - Oracle base directory, 2-28
 - Oracle home directory, 2-30
 - Oracle Inventory directory, 2-29
 - oraInventory, 2-29
 - permission for data file directories, 2-34
- disc
 - mounting, 4-5
- disk devices
 - in Automatic Storage Management, 1-12
 - managing with Automatic Storage Management, 1-11
- disk space
 - checking, 2-5
 - requirement for Oracle base directory, 2-31
 - requirements for preconfigured database in Automatic Storage Management, 3-7
- disks
 - checking availability for Automatic Storage Management, 3-12
 - configuring for Automatic Storage Management, 3-6 to ??, 3-10
 - displaying attached disks, 3-12
 - supported for Automatic Storage Management, 3-10
- DISPLAY environment variable
 - setting, 2-38
- DOMAIN_NAME initialization parameter, 6-10
- Dynamic Host Configuration Protocol. *See* DHCP

E

- Enterprise Manager Database Control Agent
 - ports
 - ranges and protocol, E-2
- enterprise.rsp file, A-5
- environment
 - configuring for oracle user, 2-38
- environment variables
 - DISPLAY, 2-38
 - NLS_LANG, F-3
 - ORACLE_BASE, 2-32, 2-38
 - ORACLE_HOME, 2-37, 2-38, 2-40
 - ORACLE_HOSTNAME, 2-17
 - ORACLE_SID, 2-38
 - PATH, 2-38
 - SHELL, 2-38
 - TMP and TMPDIR, 2-5, 2-39, 3-3
 - TNS_ADMIN, 2-40
- errata
 - Linux kernel errata, 2-8
- errors
 - configuration assistants, G-4

- display errors, G-2
- installation, G-3, G-5
- remote terminal installation, G-2
- response file installation, G-5
- silent mode, G-5
- su command, G-2
- /usr/X11R6/bin/xdpyinfo, G-2
- X Window, G-1
- X Window display errors, G-2
- /etc/csh.login file, 2-24
- /etc/csh.login.local file, 2-24
- /etc/pam.d/login file, 2-24
- /etc/profile file, 2-24
- /etc/profile.local file, 2-24
- /etc/security/limits.so file, 2-24
- /etc/sysctl.conf file, 2-26
- EXAMPLE tablespaces
 - description, 6-12
 - example01.DBF data file, 6-12
- example01.DBF data file, 6-12
- examples
 - Automatic Storage Management failure groups, 3-8
 - Oracle base directories, 2-29
- external redundancy
 - Automatic Storage Management redundancy level, 3-7

F

- failure group
 - examples of Automatic Storage Management failure groups, 3-8
- failure groups
 - characteristics of Automatic Storage Management failure group, 3-8
 - examples in Automatic Storage Management, 3-8
 - in Automatic Storage Management, 1-12
- Fast Recovery Area, 5-5
- fatal errors, G-5
- fdisk command, 3-12
- file mode creation mask
 - setting, 2-38
- file sets, 2-6
- file system
 - appropriate for Oracle base directory, 2-32
 - data file and recovery file placement options, 2-33
 - requirements for Oracle base directory, 2-32
 - using for data files, 2-33
- file-max file, 2-25
- file-max parameter
 - recommended value on Linux x86, 2-25
- files, D-4
 - \$ORACLE_HOME/dbs/initsid.ora, 6-11
 - \$ORACLE_HOME/install/portlist.ini, 6-2
 - control, 6-13, D-4
 - data files, D-4
 - dbca.rsp, A-5
 - default shell startup file, 2-24

- enterprise.rsp, A-5
- /etc/csh.login, 2-24
- /etc/csh.login.local, 2-24
- /etc/group, D-3
- /etc/pam.d/login, 2-24
- /etc/passwd, D-3
- /etc/profile, 2-24
- /etc/profile.local, 2-24
- /etc/security/limits.so, 2-24
- /etc/sysctl.conf, 2-26
- listener.ora, 5-7
- mgw.ora, 5-9
- oraInst.loc, 2-21
- oraInst.loc file, A-3
- oratab, 2-30
- /proc/sys/fs/file-max, 2-25
- /proc/sys/kernel/sem, 2-25
- /proc/sys/kernel/shmall, 2-25
- /proc/sys/kernel/shmmax, 2-25
 - shmmax file, 2-25
- /proc/sys/kernel/shmmni, 2-25
- /proc/sys/net/ipv4/ip_local_port_range, 2-25
- redo log, 6-12
- response files, A-4
- tnsnames.ora, 5-6
- Flash Recovery Area
 - See Fast Recovery Area
- For, 2-32
- free
 - UNIX command, 2-4, 3-2

G

- Gateways products FAQ, H-9
- Global Database Name
 - about, 4-12
- global database name, 6-10
- globalization support, F-1
- Grid Control. *See* Oracle Enterprise Manager Grid Control
- group file, D-3
- groups
 - checking for existing oinstall group, 2-21
 - creating the asm group, 2-22
 - creating the asmdba group, 2-22
 - creating the dba group, 2-21
 - creating the oinstall group, 2-20
 - creating the oper group, 2-22
 - UNIX OSDBA group (dba), 2-19
 - UNIX OSDBA group for Oracle Restart (dba), 2-19
 - UNIX OSOPER group (oper), 2-19

H

- hardware certification, 1-5
- hardware requirements, 2-3, 3-2
 - disk space, 2-5
 - display, 2-6
 - memory, 2-3

- run level, 2-6
- system architecture, 2-4
- high redundancy
 - Automatic Storage Management redundancy level, 3-7
- home directory
 - See Oracle home directory
- host name, setting before installation, 2-17

I

- IBM DB2 database, H-9
- IBM DB2 z/OS database, H-9
- IBM DB2/400 database, H-9
- IBM WebSphere MQ
 - requirement on Linux, 2-15
- IBM WebSphere MQ Series databases, H-9
- IDE disks
 - device names, 3-12
- Informix Server database, H-10
- initialization parameter file
 - description, 6-11
 - in databases, 6-11
 - initsid.ora, 6-11
- initialization parameters
 - DB_NAME, 6-10
 - DOMAIN_NAME, 6-10
 - SERVICE_NAMES, 6-10
- initsid.ora file, 6-11
- initsid.ora initialization parameter file, 6-11
- installation
 - accessing installation software, 4-3
 - Automatic Storage Management requirements, 3-7
 - available products, 1-8
 - cleaning up after a failed installation, G-6
 - clusters, installation guidelines, 4-2
 - component-specific guidelines, 4-1
 - computer aliases, multiple, 2-17
 - considerations, 1-4
 - database editions, 1-8
 - errors, G-3, G-5
 - silent mode, G-5
 - laptops, 2-17
 - log files, G-3
 - overview, 1-1 to 1-16
 - response file
 - oraInst.loc file, A-3
 - response files, A-1, A-4
 - preparing, A-4, A-5
 - silent mode, G-5
 - templates, A-4
 - responsefile
 - error handling, G-6
 - silent mode, A-7
 - upgrading, H-3
- installation errors
 - steps to resolve, G-3
- installation guidelines, 4-9
- installation overview, 1-1

- installation software
 - copying to a hard disk, 4-6
 - extracting, 4-5
- installation software, accessing, 4-3
- Installing
 - Oracle restart, 3-14
- installing Linux, 1-2
 - default install, 1-3
 - minimal install, 1-3
- instance
 - instance identifier (SID), 2-38
- IP addresses, multiple, 2-17
- ip_local_port_range file, 2-25
- ip_local_port_range parameter
 - recommended value on Linux x86, 2-25
- IPC protocol address
 - Oracle Messaging Gateway setting, 5-8

J

- JDK
 - internationalization class, 5-9
 - run-time class, 5-9
- JDK requirements, 2-6

K

- Kernel
 - requirements, 2-7
- kernel
 - Linux errata, 2-8
- kernel parameters
 - changing, 2-26
 - configuring, 2-25

L

- languages
 - installing Oracle components in different languages, F-4
 - using Oracle components in different languages, F-3
- laptops, installing Oracle Database on, 2-17
- limits.so file, 2-24
- Linux
 - kernel errata, 2-8
 - setting shell limits, 2-24
- listener
 - identifying Oracle home for, 2-37
 - lsnrctl command, 2-38
 - stopping, 2-37, 2-38
- listener.ora file, 5-7
 - modifying for external procedures, 5-8
- local device
 - using for data files, 2-33
- log files, G-3
 - troubleshooting, G-3
- logical volume manager
 - See* LVM
- login file, 2-24
- loopback adapters

- non-networked computers, 2-17
- lsdev command, 3-12
- lsnrctl command, 2-38
- LVM
 - recommendations for Automatic Storage Management, 3-7

M

- mask
 - setting default file mode creation mask, 2-38
- memory requirements, 2-3, 3-2
- MGW_AGENT service name, 5-8
- mgwextproc service
 - adding static service information, 5-8
- mgw.ora file
 - modifying, 5-9
- Microsoft SQL Server database, H-10
- migrating
 - See* upgrading
- migrating applications to Oracle, H-8
- migrating non-Oracle databases to Oracle, H-3
- mirroring Automatic Storage Management disk groups, 3-7
- mkdir command, 2-32, 2-34
- mode
 - setting default file mode creation mask, 2-38
- mount point
 - for Oracle base directory, 2-28
- mount point directories, 4-7
- mount point directory
 - choosing, C-2
- mount points
 - Optimal Flexible Architecture conventions for creating, D-1
- MQSeries
 - class, 5-9
- multihomed computers, installing on, 2-17
- multiple aliases, computers with, 2-17
- multiple databases and ASM, 2-20
- multiple Oracle homes, 1-5

N

- naming subdirectories, D-3
- NAS devices
 - creating files on for use with Automatic Storage Management, C-4
 - guidelines for configuration, C-1
- Net Configuration Assistant
 - troubleshooting, G-4
- Net Configuration Assistant (NetCA)
 - response files, A-8
 - running at command prompt, A-8
- netca.rsp file, A-5
- network adapters
 - computers with multiple aliases, 2-17
 - non-networked computers, 2-17
 - primary, on computers with multiple aliases, 2-17

- See also* loopback adapters, primary network adapters
- network cards, multiple, 2-17
- Network Information Services
 - alternative to local users and groups, 2-20
- Network Information Services.*See* NIS
- network setup
 - about, 2-16
 - computers with multiple aliases, 2-17
- network topics
 - DHCP computers, 2-16
 - laptops, 2-17
 - multiple network cards, 2-17
 - non-networked computers, 2-17
- NFS
 - mount options, C-5
- NLS_LANG environment variable, F-3
- nofile
 - shell limit on Linux x86, 2-24
- noninteractive mode
 - See also* response files, response file mode, A-2
- non-networked computers, 2-17
- non-Oracle databases, listed, H-9
- normal redundancy, Automatic Storage Management
 - redundancy level, 3-7
- nproc
 - shell limit on Linux x86, 2-24

O

- OEM
 - See* Oracle Enterprise Manager
- oinstall group
 - checking for existing, 2-21
- oinstall groups
 - creating, 2-20
- OLAP tools
 - about, H-6
 - Oracle OLAP, H-6
- OMF
 - See* Oracle Managed Files
- oper group
 - creating, 2-22
 - description, 2-19
- oper groups
 - creating, 2-22
- operating system
 - checking distribution and version, 2-7
- operating system accounts
 - creating and configuring, 5-3
- operating system groups
 - creating the oinstall group, 2-20
- operating system requirements, 2-6
- operating system users
 - root user, 4-9
- Optimal Flexible Architecture
 - Automatic Storage Management, D-5
 - conventions for creating mount points, D-1
 - file identification, D-6
 - file mapping, D-6
 - files systems, D-2
 - naming, D-1
 - database files, D-4
 - Oracle base directory, D-2
 - subdirectories, D-3
 - very large databases, D-3
 - Oracle Managed Files, D-5
 - pathnames, D-3
 - recommendations for Oracle base directory, 2-28
 - recommended path for Oracle base directory, 2-28
 - recommended path for Oracle home directory, 2-30
 - recommended path for Oracle Inventory directory, 2-29
 - special tablespaces, D-5
 - standard, D-1
 - using separate segments, D-5
- Oracle ACFS, 1-11, 3-3
 - requirements, 3-3
- Oracle ADVM, 3-3
 - requirements, 3-3
- Oracle Application Server, H-4
- Oracle applications
 - installing with Oracle Database, H-8
- Oracle Automatic Storage Management Cluster File System, 1-11, 3-3
- Oracle base directory
 - creating, 2-32
 - creating new, 2-31
 - description, 2-28
 - determining disk space on, 2-31
 - disk space requirements, 2-31
 - equivalent directory on Microsoft Windows, 2-28
 - examples, 2-29
 - identifying appropriate file system, 2-32
 - identifying existing, 2-30
 - mount point for, 2-28
 - naming conventions, D-2
 - recommended path, 2-28
 - relationship with Oracle software owner user, 2-28
 - requirement for, 2-28
 - requirements for existing directory, 2-31
 - requirements on file system, 2-32
- Oracle Cluster Registry
 - See* OCR
- Oracle Cluster Registry port, E-3
- Oracle Clusterware
 - about, H-3
 - ports, E-3
 - ports, ranges and protocol, E-3
 - used with Oracle Real Application Clusters, H-3
- Oracle components
 - using in different languages, F-3
- Oracle Data Mining
 - about, H-7
 - installing, H-7
- Oracle Database
 - administering and monitoring, H-4

- creating data file directories, 2-34
- Enterprise Edition installation, 1-8
- getting started using
 - accessing, 6-4
 - starting and stopping database, 6-4
- installing with Oracle applications, H-8
- minimum disk space requirements, 2-33
- naming, 4-12
- requirements with Automatic Storage Management, 3-7
- security management, H-5
- setting ORACLE_SID environment variable, 2-38
- Standard Edition installation, 1-8
- upgrading, H-3
- Web servers, H-8
- Oracle Database Client
 - configuring connections, H-2
- Oracle Database components
 - administering and monitoring, H-4
 - connectivity FAQ, H-9
 - FAQ on installing, H-2 to H-3
 - installing with Oracle applications, H-8
 - installing with Oracle Database tools, H-4
- Oracle Database Configuration Assistant
 - response file, A-5
- Oracle Database Recovery Manager (RMAN)
 - about, H-7
- Oracle Database SID
 - about, 4-12
 - naming rules, 4-12
- Oracle Database Vault
 - audit policy, 1-7
 - postinstallation task, 5-7
 - preinstallation requirement, 2-15
- Oracle Enterprise Management Agent
 - HTTP port, changing, E-3
- Oracle Enterprise Manager, 1-13
 - Database Control
 - logging into, 6-1
 - port number, 6-1
 - using to modify control files, 6-13
 - using to modify redo log files, 6-13
 - using to view control files, 6-13
 - using to view redo log files, 6-13
 - login privileges, 6-2
- Oracle Enterprise Manager Database Control
 - ports, changing, E-3
 - ports, ranges and protocol, E-2
- Oracle Event Manager
 - ports, ranges and protocol, E-3
- Oracle home
 - cloning, B-1
- Oracle home directory
 - description, 2-30
 - identifying for listener, 2-37
 - multiple homes, network considerations, 2-17
 - recommended path, 2-30
 - requirement for, 2-30
 - requirements, 2-30
 - using to identify Oracle base directory, 2-30
- Oracle home name, 2-30
- Oracle host name, setting before installation, 2-17
- Oracle internationalization class, 5-9
- Oracle Internet Directory, H-5
- Oracle Inventory
 - description, 2-29
 - pointer file, 2-21
- Oracle Inventory directory
 - description, 2-29
 - recommended path, 2-29
- Oracle Inventory group
 - creating, 2-21
- Oracle Inventory groups
 - checking for existing, 2-21
 - creating, 2-20
- Oracle JDBC class, 5-9
- Oracle Label Security
 - post-installation tasks, 5-7
- Oracle Managed Files
 - Optimal Flexible Architecture naming conventions, D-5
- Oracle Messaging Gateway
 - CSD requirements
 - on Linux, 2-15
 - postinstallation tasks, 5-7
 - requirements on Linux, 2-15
- Oracle Messaging Gateway class, 5-9
- Oracle Migration Workbench
 - migrating non-Oracle applications to Oracle, H-8
 - migrating non-Oracle databases to Oracle, H-3
- Oracle Net
 - configuration file directory, 5-6
 - identifying Oracle home for listener, 2-37
 - lsnrctl command, 2-38
 - stopping listener, 2-37
 - stopping the listener, 2-38
- Oracle Net Configuration Assistant
 - response file, A-5
- Oracle Net Services
 - post-installation tasks, 5-6
- Oracle OLAP
 - about, H-6
- Oracle Precompilers
 - postinstallation tasks, 5-9
- Oracle Procedural Gateway
 - listed products, H-9
- Oracle Real Application Clusters (RAC)
 - installed before Oracle Database, 4-2
 - installing with Oracle Enterprise Manager, H-5
 - Oracle Clusterware
 - about, H-3
- Oracle Restart
 - description, 3-1
 - Installing, 3-14
 - OSDBA group description, 2-19
 - user, 2-19
- Oracle Schemas, xiv
- Oracle software
 - removing, 7-5
- Oracle Software Owner user

- creating, 2-22
- oracle user, 2-23
- Oracle software owner user
 - configuring environment for, 2-38
 - determining default shell, 2-38
 - relationship with Oracle base directory, 2-28
- Oracle SQL Developer
 - accessing, 6-4
- Oracle SQL*Net Listener
 - ports, ranges and protocol, E-2
- Oracle Technology Network (OTN)
 - downloading documentation from, xiv
- Oracle Text knowledge base, 5-10
- Oracle Transparent Gateway
 - listed products, H-9
- Oracle Universal Installer
 - guidelines for using, 4-1
 - installation guidelines, 4-1
 - response files, A-1
 - list of, A-5
 - running, 4-8
 - running in different languages, F-4
- oracle user
 - configuring environment for, 2-38
 - creating, 2-22
 - determining default shell, 2-38
 - relationship with Oracle base directory, 2-28
- Oracle user accounts
 - configuring, 5-3
- Oracle Validated Configuration RPM
 - about, 1-3
 - installing, 1-4
- Oracle XML DB
 - about, H-6
 - ports, E-4
 - ports, changing, E-4
 - ports, ranges and protocol, E-3
- ORACLE_BASE environment variable, 2-32
 - setting, 2-38
- ORACLE_HOME environment variable
 - setting, 2-37
 - unsetting, 2-40
- ORACLE_HOSTNAME, 2-17
- ORACLE_HOSTNAME environment variable
 - computers with multiple aliases, 2-17
 - multihomed computers, 2-17
 - setting before installation, 2-17
- ORACLE_SID environment variable
 - setting, 2-38
- oraInst.loc file
 - location, 2-21
 - location of, 2-21
- oraInventory directory
 - See* Oracle Inventory directory
- oratab file, 2-30
 - formats, 2-30
 - location of, 2-30
- OSASM groups
 - creating, 2-22
 - multiple databases, 2-20

- SYSASM, 2-20
- OSDBA groups
 - creating, 2-21
 - creating for Oracle Grid Infrastructure, 2-22
 - description for database, 2-19
 - SYSDBA privilege, 2-19
 - SYSDBA privilege for Oracle Restart, 2-19
- OSOPER groups
 - creating, 2-22
 - description for database, 2-19
 - SYSOPER privilege, 2-19
- OTN Web site
 - downloading installation software from, 4-3

P

- package requirements, 2-8
 - Linux x86, 2-9
 - Linux x86-64, 2-11
- packages, checking, 2-13
- partition
 - using with Automatic Storage Management, 3-7
- partitions
 - creation for Automatic Storage Management
 - disks, 3-10
- passwd file, D-3
- password file for Automatic Storage Management, 3-5
- passwords
 - resetting, 6-8
 - with Database Control, 6-8
 - with SQL*Plus, 6-9
 - reviewing, 6-5
 - specifying for response files, A-2
 - unlocking, 6-8
 - with Database Control, 6-8
 - with SQL*Plus, 6-9
 - See also* security
- PATH environment variable
 - setting, 2-38
- pathnames
 - Optimal Flexible Architecture, D-3
- permissions
 - for data file directories, 2-34
 - for Oracle base directory, 2-32
- port numbers
 - managing, E-1
- portlist.ini file, 6-2, E-2
- ports
 - access URLs, E-2
 - Cluster Manager, ranges and protocol, E-3
 - Cluster Synchronization Services, ranges and protocol, E-3
 - configured for applications, E-2
 - Connection Manager, ranges and protocol, E-2
 - default ranges, E-1
 - Enterprise Manager Database Control Agent, ranges and protocol, E-2
 - Oracle Cluster Registry, E-3
 - Oracle Clusterware, E-3

- Oracle Clusterware, ranges and protocol, E-3
- Oracle Enterprise Management Agent HTTP, changing, E-3
- Oracle Enterprise Manager Database Control, changing, E-3
- Oracle Enterprise Manager Database Control, ranges and protocol, E-2
- Oracle Event Manager, ranges and protocol, E-3
- Oracle Net, E-2
- Oracle SQL*Net Listener, ranges and protocol, E-2
- Oracle XML DB, E-4
- Oracle XML DB, changing, E-4
- Oracle XML DB, ranges and protocol, E-3
- post-installation
 - required tasks, 5-1
 - Oracle Label Security, configuring, 5-7
 - Oracle Net Services, configuring, 5-6
 - patches, installing and downloading, 5-1
- postinstallation
 - recommended tasks
 - client static library, generating, 5-4
 - creating operating system accounts, 5-3
 - root.sh script, backing up, 5-2
 - required tasks
 - configuring Oracle Messaging Gateway, 5-7
 - Oracle Precompilers, 5-9
- postinstallation tasks
 - Oracle Text knowledge base, 5-10
- preconfigured database
 - Automatic Storage Management disk space requirements, 3-7
 - requirements when using Automatic Storage Management, 3-7
- Pro*C/C++
 - configuring, 5-9
 - See also C compiler
- process
 - stopping existing, 2-37
 - stopping listener process, 2-37
- /proc/sys/fs/file-max file, 2-25
- /proc/sys/kernel/sem file, 2-25
- /proc/sys/kernel/shmall file, 2-25
- /proc/sys/kernel/shmmni file, 2-25
- /proc/sys/net/core/rmem_default file, 2-25
- /proc/sys/net/core/rmem_max file, 2-25
- /proc/sys/net/core/wmem_default file, 2-25
- /proc/sys/net/core/wmem_max file, 2-25
- /proc/sys/net/ipv4/ip_local_port_range file, 2-25
- profile file, 2-24
- profile.local file, 2-24
- ps command, 2-37

R

- RAID
 - device names, 3-12
 - using for Oracle data files, 2-33
- RAM requirements, 2-3, 3-2
- readme.txt file, E-2

- reconfiguring CSS, 7-1
- recovery files
 - options for placing on file system, 2-33
- recovery of databases
 - Oracle Backup and Recovery, H-7
- Red Hat Package Manager
 - See RPM
- redo log, D-4
- redo log files
 - in starter database, 6-12
 - locating, 6-12
 - naming, D-4
 - reviewing, 6-11
 - using Oracle Enterprise Manager Database Control with, 6-13
- redundancy level
 - and space requirements for preconfigured database, 3-7
 - for Automatic Storage Management, 3-7
- redundant array of independent disks
 - See RAID
- Removing Oracle software, 7-5
- Rendezvous
 - requirement on Linux, 2-15
- requirements
 - hardware, 2-3, 3-2
- response file installation
 - oraInst.loc file, A-3
 - response files
 - preparing, A-4, A-5
 - templates, A-4
 - silent mode, A-7
 - errors, G-5
- response file mode
 - about, A-2
 - reasons for using, A-3
- response files, A-1
 - about, A-1
 - Automatic Storage Management, A-3
 - creating with template, A-4
 - dbca.rsp, A-5
 - enterprise.rsp, A-5
 - general procedure, A-3
 - Net Configuration Assistant, A-8
 - netca.rsp, A-5
 - passing values at command line, A-2
 - passwords, A-2
 - security, A-2
 - specifying with Oracle Universal Installer, A-7
- response files installation
 - about, A-1
- rmem_default file, 2-25
- rmem_default parameter
 - recommended value on Linux, 2-25
- rmem_max file, 2-25
- rmem_max parameter
 - recommended value on Linux, 2-25
- root user, 4-9
 - logging in as, 2-2
- root.sh script

- backing up, 5-2
- RPM
 - checking, 2-13
 - rpm command, 2-13

S

- Sample Schemas
 - tablespaces and data files, 6-12
- SAN (storage area network) disks, 3-10
- schema passwords, 4-14
- schemas
 - database schema passwords, 4-14
 - Oracle Schemas, about, xiv
 - Sample Schemas tablespaces and data files, 6-12
- SCSI disks
 - device names, 3-12
- SE Linux, 2-7, 4-7
- security
 - dividing ownership of Oracle software, 2-18
 - management tools, H-5
 - See also* passwords
- Security Enhanced Linux, 2-7, 4-7
- sem file, 2-25
- semnmi parameter
 - recommended value on Linux x86, 2-25
- semnms parameter
 - recommended value on Linux x86, 2-25
- semmsl parameter
 - recommended value on Linux x86, 2-25
- semopm parameter
 - recommended value on Linux x86, 2-25
- server parameter file (SPFILE), 3-5
- SERVICE_NAMES initialization parameter, 6-10
- shell
 - default shell startup file, 2-24
 - determining default shell for oracle user, 2-38
- SHELL environment variable
 - checking value of, 2-38
- shell limits
 - setting on Linux x86, 2-24
- shmall file, 2-25
- shmall parameter
 - recommended value on Linux x86, 2-25
- shmmax parameter
 - recommended value on Linux x86, 2-25
- shmmni file, 2-25
- shmmni parameter
 - recommended value on Linux x86, 2-25
- SID, 6-10
 - setting ORACLE_SID environment variable, 2-38
- SID. *See* Oracle Database SID
- silent mode
 - about, A-1
 - reasons for using, A-2
 - See also* response file mode, response files, A-1
- silent mode installation, A-7
- software certification, 1-5
- software requirements, 2-6
- SPFILE server parameter file, 3-5

- SQL Developer
 - accessing, 6-4
- SQL Server database, H-10
- SQL*Plus
 - accessing, 6-4
- SQLJ class, 5-9
- startup file
 - default shell startup file, 2-24
- static service information
 - adding for mgwextproc service, 5-8
- storage area network disks, 3-10
- storage devices
 - configuring for datafiles, 2-35
- storage management *See* Automatic Storage Management
- suppressed mode. *See* response file mode
- swap space
 - checking, 2-4, 3-2
 - requirements, 2-3, 3-2
- Sybase Adapter Server database, H-10
- SYSASM
 - OSASM, 2-20
- sysctl command, 2-26
- sysctl.conf file, 2-26
- SYSDBA privilege
 - associated UNIX group, 2-19
- SYSOPER privilege
 - associated UNIX group, 2-19
- SYSTEM
 - tablespace, description, 6-12
- System Identifier, 6-10
 - See* SID
- system01.dbf data file, 6-12

T

- tablespaces, 6-12
 - defined, 6-11
 - in databases, 6-11
 - reviewing, 6-11
 - setting up, 6-11
 - expanding for large sorts, 6-12
 - Optimal Flexible Architecture
 - special tablespaces, D-5
 - SYSTEM, 6-12
 - TEMP, 6-12
 - UNDOTBS, 6-12
 - USERS, 6-12
- TEMP
 - tablespace (temp01.dbf), 6-12
- temp01.dbf data file, 6-12
- temporary disk space
 - requirements, 2-3, 3-2
- Teradata database, H-10
- TIBCO Rendezvous
 - requirement on Linux, 2-15
- TMP environment variable, 2-5, 3-3
 - setting, 2-39
- TMPDIR environment variable, 2-5, 3-3
 - setting, 2-39

- TNS_ADMIN environment variable
 - unsetting, 2-40
- tnsnames.ora file, 5-6
 - adding a connect descriptor, 5-8
 - MGW_AGENT service name, 5-8
 - modifying for external procedures, 5-8
- troubleshooting, G-1
 - display errors, G-2
 - fatal errors, G-5
 - remote terminal installation, G-2
 - su command, G-2
 - /usr/X11R6/bin/xdpyinfo, G-2

U

- ulimit command, 2-24
- umask command, 2-38
- UNDOTBS
 - tablespace (undotbs01.dbf), 6-12
- UNIX commands
 - chmod, 2-32, 2-34
 - chown, 2-32, 2-34
 - fdisk, 3-12
 - free, 2-4, 3-2
 - lsdev, 3-12
 - mkdir, 2-32, 2-34
 - ps, 2-37
 - rpm, 2-13
 - sysctl, 2-26
 - ulimit, 2-24
 - umask, 2-38
 - unset, 2-40
 - unsetenv, 2-40
 - xhost, 2-2
 - xterm, 2-2
- UNIX groups
 - checking for existing oinstall group, 2-21
 - OSDBA (dba), 2-19
 - OSDBA (dba) for Oracle Restart, 2-19
 - OSOPER (oper), 2-19
 - using NIS, 2-20
- UNIX users
 - using NIS, 2-20
- UNIX workstation
 - installing from, 2-2
- unset command, 2-40
- unsetenv command, 2-40
- upgraded databases
 - configuring, 5-3
- upgrading, 1-16
 - ASM, 1-17
 - Daylight Savings Time, 1-17
 - Oracle Database on RHEL 2.1, 1-16
- useradd command, 2-23
- USERS
 - tablespace (users01.dbf), 6-12
- users
 - creating the oracle user, 2-22
 - Oracle Restart, 2-19
- users and groups, 2-18

- UTLRP.SQL
 - recompiling invalid SQL modules, 5-3

V

- very large databases
 - Optimal Flexible Architecture naming mount points, D-3

W

- Web servers (Oracle HTTP Server), H-8
- WebSphere MQ
 - CSD download location, 2-15
 - CSDs required
 - on Linux, 2-15
 - requirement on Linux, 2-15
- WebSphere MQ class, 5-9
- WebSphere MQ Series database, H-9
- Windows
 - analogy for Oracle base directory, 2-28
- wmem_default file, 2-25
- wmem_default parameter
 - recommended value on Linux, 2-25
- wmem_max file, 2-25
- wmem_max parameter
 - recommended value on Linux, 2-25

X

- X Window
 - display errors, G-1
- X Window system
 - enabling remote hosts, 2-2
- xhost command, 2-2
- XML data, H-6
- xterm command, 2-2

