

# Oracle8i

Administrator's Reference

Release 3 (8.1.7) 64 Bit for SGI IRIX

May 2001

Part No. A87435-01

**ORACLE®**

Part No. A87435-01

Copyright © 1996, 2001, Oracle Corporation. All rights reserved.

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

The programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the programs is prohibited.

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

If the programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

**Restricted Rights Notice** Programs delivered subject to the DOD FAR Supplement are "commercial computer software" and use, duplication, and disclosure of the programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, programs delivered subject to the Federal Acquisition Regulations are "restricted computer software" and use, duplication, and disclosure of the programs shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software - Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

Oracle is a registered trademark, and ConText, *interMedia*, JDeveloper, Net8, Oracle Call Interface, Oracle Database Configuration Assistant, Oracle Data Migration Assistant, Oracle Developer, Oracle Enterprise Manager, Oracle Names, Oracle Spatial, Oracle Time Series, Oracle Video Information Retrieval, Oracle7, Oracle8i, PL/SQL, Pro\*C/C++, Pro\*COBOL, Pro\*FORTRAN, SQL\*Loader, and SQL\*Plus are trademarks or registered trademarks of Oracle Corporation. Other names may be trademarks of their respective owners.

---

---

# Contents

<b>Send Us Your Comments .....</b>	<b>ix</b>
<b>Preface.....</b>	<b>xi</b>
Purpose of this Guide.....	xi
Audience .....	xi
Oracle8i and Oracle8i Enterprise Edition.....	xi
Typographic Conventions.....	xi
Command Syntax .....	xii
Accessing Installed Documentation.....	xiii
Related Documentation .....	xiv
Oracle Services and Support .....	xiv
<b>1 Administering Oracle8i</b>	
<b>Overview .....</b>	<b>1-2</b>
<b>Environment Variables .....</b>	<b>1-2</b>
Oracle8i Environment Variables .....	1-3
UNIX Environment Variables.....	1-5
Setting a Common Environment.....	1-8
Setting and Exporting the Value of a Variable in a Current Session .....	1-9
Setting the System Time .....	1-9
<b>Product Executables.....</b>	<b>1-9</b>
<b>System Global Area.....</b>	<b>1-10</b>
<b>Calculating the Size of the SGA .....</b>	<b>1-11</b>
<b>Oracle8i Memory Requirements and Usage.....</b>	<b>1-13</b>

<b>Server Resource Limits</b> .....	1-14
<b>Database Limits</b> .....	1-14
<b>Special Accounts and Groups</b> .....	1-15
<b>Security</b> .....	1-16
Groups and Security.....	1-17
Security for Server Manager Commands.....	1-17
Security for Database Files .....	1-17
Remote Passwords.....	1-18
Customizing the initsid.ora File .....	1-19
<b>Embedded PL/SQL Gateway</b> .....	1-21
Overview.....	1-21
Installing the Embedded PL/SQL Gateway.....	1-22
<b>Demonstrations Files</b> .....	1-23
SQL*Loader Demonstrations .....	1-23
Administering SQL*Loader.....	1-24
PL/SQL Demonstrations.....	1-26
Database Examples.....	1-28

## 2 Tuning Oracle8i

<b>The Importance of Tuning</b> .....	2-2
<b>SGI IRIX Tools</b> .....	2-2
sar .....	2-2
swap.....	2-3
<b>SQL Scripts</b> .....	2-3
<b>Tuning Memory Management</b> .....	2-3
Allocate Sufficient Swap Space .....	2-4
Control Paging .....	2-4
Hold the SGA in a Single Shared Memory Segment .....	2-5
<b>Tuning Disk I/O</b> .....	2-5
Tune the Database Writer to Increase Write Bandwidth .....	2-5
Choose the Appropriate File System Type .....	2-7
<b>Monitoring Disk Performance</b> .....	2-7
Disk Performance Issues .....	2-7

<b>Tuning CPU Usage</b> .....	2-8
Keep All Oracle Users and Processes at the Same Priority .....	2-8
Use Processor Affinity and Binding on Multi-Processor Systems .....	2-8
Use a Client/Server Configuration.....	2-8
Use the Post-Wait Driver.....	2-8
Use Single-Task Linking for Large Exports/Imports and SQL*Loader Jobs .....	2-9
<b>Tuning Oracle Resource Contention and UNIX Kernel Parameters</b> .....	2-9
<b>Tuning and Specifying Block Size and File Size</b> .....	2-10
<b>Tuning the SGI IRIX Buffer Cache Size</b> .....	2-11
<b>Tuning Oracle8i on IRIX Systems</b> .....	2-11
Lock the SGA in Physical Memory .....	2-14
<b>Using Trace and Alert Files</b> .....	2-15
Trace File Names .....	2-15
Alert Files.....	2-15
<b>Using Raw Devices/Volumes</b> .....	2-16
Disadvantages of Raw Devices/Volumes .....	2-16
Guidelines for Using Raw Devices/Volumes .....	2-17
Raw Device Setup.....	2-18

### 3 Administering SQL\*Plus

<b>Administering SQL*Plus</b> .....	3-2
Using Setup Files .....	3-2
Using the PRODUCT_USER_PROFILE Table.....	3-3
Using Demonstration Tables .....	3-3
Installing and Removing the Help Facility.....	3-4
<b>Using SQL*Plus</b> .....	3-5
Using a System Editor from SQL*Plus.....	3-5
Running Operating System Commands from SQL*Plus.....	3-6
Interrupting SQL*Plus .....	3-6
Using the SPOOL Command.....	3-7
<b>Restrictions of SQL*Plus</b> .....	3-7
Resizing Windows.....	3-7
Return Codes.....	3-7

<b>4</b>	<b>Using Oracle Precompilers and the Oracle Call Interface</b>	
	<b>Overview of Oracle Precompilers</b> .....	4-2
	Precompiler Configuration Files .....	4-2
	Issues Common to All Precompilers.....	4-3
	Static and Dynamic Linking.....	4-3
	Client Shared Library .....	4-4
	<b>Pro*C/C++ Precompiler</b> .....	4-4
	Pro*C/C++ Demonstration Programs.....	4-5
	Pro*C/C++ User Programs .....	4-6
	<b>Pro*FORTRAN Precompiler</b> .....	4-7
	Pro*FORTRAN Demonstration Programs.....	4-7
	Pro*FORTRAN User Programs .....	4-8
	<b>Oracle Call Interface</b> .....	4-9
	OCI Demonstration Programs .....	4-9
	OCI User Programs .....	4-10
	<b>Custom Make Files</b> .....	4-11
	<b>Correcting Undefined Symbols</b> .....	4-11
	<b>Multi- threaded Applications</b> .....	4-12
	Using Signal Handlers .....	4-12
	<b>XA Functionality</b> .....	4-15
<b>5</b>	<b>Configuring Net8</b>	
	<b>Core Net8 Products and Features</b> .....	5-2
	Net8 Files and Utilities.....	5-2
	Oracle Connection Manager.....	5-3
	Server Models.....	5-4
	Oracle Names .....	5-5
	<b>Net8 Protocol Support</b> .....	5-5
	ADDRESS Specification.....	5-5
	<b>BEQ Protocol</b> .....	5-6
	<b>IPC Protocol</b> .....	5-8
	<b>RAW Protocol</b> .....	5-9
	<b>TCP/IP Protocol Adapter</b> .....	5-10
	Specifying a TCP/IP ADDRESS .....	5-10

<b>Configuring Oracle Intelligent Agent for Oracle SNMP</b> .....	5-11
<b>Oracle Advanced Security</b> .....	5-13

## **6 Running Oracle Data Option Demonstrations**

<b>Oracle8i interMedia</b> .....	6-2
Text .....	6-2
Audio, Video, and Image .....	6-2
Locator.....	6-6
Web Agent and Clipboard .....	6-8
<b>Oracle8i Time Series Demonstrations</b> .....	6-8
<b>Oracle8i Visual Information Retrieval</b> .....	6-9
<b>Oracle8i Spatial</b> .....	6-11

## **A Optimal Flexible Architecture**

<b>Optimal Flexible Architecture</b> .....	A-2
Characteristics of an OFA-Compliant Database .....	A-2
<b>Optimal Flexible Architecture Implemented on UNIX</b> .....	A-4
Mount Points .....	A-4
Naming Directories .....	A-5
Naming Database Files .....	A-8
Separate Segments with Different Requirements .....	A-9
Naming Tablespaces .....	A-10
Exploiting the OFA Structure for Oracle Files .....	A-10
OFA File Mapping.....	A-11
Raw Device Sizes .....	A-12
Directory Structure .....	A-13

## **Index**



---

---

# Send Us Your Comments

**Oracle8i Administrator's Reference Release 3 (8.1.7) 64 Bit for SGI IRIX**

**Part No. A87435-01**

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

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

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

- Email - [osdwrite\\_us@oracle.com](mailto:osdwrite_us@oracle.com)
- FAX - 650.413.0177 Attn: Corinne Arne
- Postal service:  
Corinne Arne  
Oracle Corporation  
500 Oracle Parkway, Mailstop 501ip4  
Redwood Shores, CA 94065  
USA

If you would like a reply, please provide your name, address, and telephone number.

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



---

---

# Preface

## Purpose of this Guide

This guide, with the *Oracle8i Release Notes Release 3 (8.1.7) 64 Bit for SGI IRIX* and the *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX*, describes how to install and configure Oracle8i release 3 (8.1.7) on SGI IRIX. For additional product information, see the *Oracle8i Generic Documentation Set*.

## Audience

This guide is intended for anyone responsible for installing Oracle8i release 3 (8.1.7) on SGI IRIX systems.

## Oracle8i and Oracle8i Enterprise Edition

Unless otherwise noted, the information in this guide is common to both Oracle8i and Oracle8i Enterprise Edition.

## Typographic Conventions

The following typographic conventions are used in this guide:

- |                        |  |
|------------------------|--|
| <code>monospace</code> | Monospace type indicates UNIX commands, directory names, user names, path names, file names, and text quoted from screen output.                                 |
| brackets [ ]           | Words enclosed in brackets indicate button and key names (for example, Press [Return]). Note that brackets have a different meaning when used in command syntax. |

<i>italics</i>	Italic type indicates variables, including variable portions of file names. It is also used for emphasis.
UPPERCASE	Uppercase letters indicate Structured Query Language (SQL) reserved words, initialization parameters, and environment variables.

## Command Syntax

All examples that show commands or file contents appear in monospace font and assume the use of the Bourne shell. The dollar prompt (\$) at the beginning of the UNIX command examples is the default UNIX command prompt. Do not enter it in the examples.

backslash \	A backslash indicates a command that is too long to fit on a single line. Enter the line as printed (with a backslash) or enter it as a single line without a backslash: <pre>dd if=/dev/rdskc0t1d0s6 of=/dev/rst0 bs=10b \ count=10000</pre>
braces { }	Braces indicate required items: <code>.DEFINE {macro1}</code>
brackets [ ]	Brackets indicate optional items: <code>cvtcrt termname [outfile]</code>  Note that brackets have a different meaning when used in regular text.
ellipses ...	Ellipses indicate an arbitrary number of similar items: <pre>CHKVAL fieldname value1 value2 ... valueN</pre>
<i>italics</i>	Italic type indicates a variable. Substitute a value for the variable: <pre>library_name</pre>
vertical line	A vertical line indicates a choice within braces or brackets: <pre>SIZE filesize [K M]</pre>

## Accessing Installed Documentation

Oracle8i release 3 (8.1.7) 64 Bit for SGI IRIX documentation includes this guide and the *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX*. You can install documentation in HTML and PDF (Adobe Portable Document Format, which requires Acrobat Reader) formats. Platform-specific documentation files are installed from the Oracle8i CD-ROM. Generic documentation files are installed from the *Oracle8i On-Line Generic Documentation CD-ROM*. The location of the installed documentation files is determined according to the following rules:

- If the ORACLE\_DOC environment variable is defined, the files are installed in that directory.
- If the ORACLE\_DOC environment variable is not defined but the ORACLE\_BASE environment variable is defined, the files are installed in the \$ORACLE\_BASE/doc directory.
- If neither the ORACLE\_DOC nor the ORACLE\_BASE environment variables are defined, the files are installed in the \$ORACLE\_HOME/doc directory.

To access the installed documentation, use a Web browser to open either the `index.htm` or `products.htm` files (the latter file does not require a frames-enabled browser). If you prefer paper documentation, you can print the PDF files.

### Oracle Product Documentation

Oracle8i product documentation is available on the *Oracle8i On-Line Generic Documentation CD-ROM*. Instructions for accessing and installing the documents on the CD-ROM are found in the README file in the top level directory of the CD-ROM.

### Oracle Information Navigator

Oracle Information Navigator is a Java-based search and navigation utility provided with Oracle online documentation. If you are using a Java-enabled browser, Information Navigator is launched automatically when you use a Web browser to open the `index.htm` file at the top level of the CD-ROM. You can use the Information Navigator with Oracle documentation, whether you are reading from the CD-ROM or from installed files.

## Related Documentation

If you are unfamiliar with the concepts or terminology associated with relational database management systems, read Chapter 1 in *Oracle8i Concepts* before beginning your installation.

Information on administering your system and tuning a production database is provided in these documents:

- *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX*
- *Oracle8i Release Notes Release 3 (8.1.7) 64 Bit for SGI IRIX*
- *Net8 Administrator's Guide*
- *Oracle8i Designing and Tuning for Performance*

Information on migrating or upgrading from a previous release of the Oracle Server is provided in *Oracle8i Migration*.

## Oracle Services and Support

Oracle Corporation offers a wide range of support services, including Oracle Education courses, Oracle Consulting services, and Oracle Worldwide Customer Support. Information on Oracle products and global services is available on the Internet, from:

<http://www.oracle.com>

The following sections provide URLs for selected services.

### Oracle Support Services

Technical Support contact information worldwide is listed at:

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

This site provides templates to help you prepare information on your problem before you call. You also need your CSI number (if applicable) or complete contact details, including any special project information relating to the problem.

### Products and Documentation

Customers in the U.S.A. can contact the Oracle Store at:

<http://store.oracle.com>

Links to stores in other countries are provided from this site.

Product documentation is available at:

<http://docs.oracle.com>

### **Customer Service**

Global Customer Service contacts are listed at:

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

### **Education and Training**

Information on training and worldwide training schedules is available at:

<http://education.oracle.com>

### **Oracle Technology Network**

Register with the Oracle Technology Network (OTN) at:

<http://technet.oracle.com>

OTN delivers technical papers, code samples, product documentation, self-service developer support, and Oracle's key developer products to enable rapid development and deployment of applications built on Oracle technology.

### **Support for the Hearing-Impaired**

To avail of TTY access to Oracle Support Services within the United States of America, phone 1-800-446-2398.



---

---

# Administering Oracle8i

This chapter describes tasks specific to administering Oracle8i on SGI IRIX. It contains the following sections:

- [Overview](#)
- [Environment Variables](#)
- [Product Executables](#)
- [System Global Area](#)
- [Calculating the Size of the SGA](#)
- [Oracle8i Memory Requirements and Usage](#)
- [Server Resource Limits](#)
- [Database Limits](#)
- [Special Accounts and Groups](#)
- [Security](#)
- [Embedded PL/SQL Gateway](#)
- [Demonstrations Files](#)

## Overview

You must set Oracle8i environment variables, parameters, memory, and user settings for Oracle8i to work. This chapter describes the various settings for Oracle8i on SGI IRIX.

In Oracle8i files and programs, a question mark (?) represents the value of the ORACLE\_HOME environment variable. For example, Oracle8i expands the question mark in the following SQL statement to the full pathname of the Oracle home directory:

```
ALTER TABLESPACE TEMP ADD DATAFILE '?/DBS/DBS2.DBF' SIZE 2M
```

The @ sign represents the ORACLE\_SID environment variable. For example, to indicate a file belonging to the current instance, enter:

```
ALTER TABLESPACE TABLESPACE_NAME ADD DATAFILE DBS2@DBF
```

## Environment Variables

This section describes the most commonly-used Oracle8i and UNIX environment variables.

To display the current value of an environment variable, use the `env` command. For example, to display the value of the ORACLE\_SID environment variable, enter:

```
$ env | grep ORACLE_SID
```

---

---

**Note:** Use the `env` command to show the value of environment variables that have been exported to the environment. Bourne shell and Korn shell can set values without exporting them.

---

---

You must define some of these variables before installing Oracle8i. The required variable settings are listed in the *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX*.

## Oracle8i Environment Variables

**Table 1-1** provides the syntax for, and examples of, environment variables used by Oracle8i.

---



---

**Note:** Do not define environment variables with values that are identical to names of Oracle Server processes, for example: arch, pmon, and dbwr.

---



---

**Table 1-1 Oracle8i Environment Variables on UNIX**

Variable	Detail	Definition
EPC_DISABLED	Function	Disables Oracle Trace.
	Syntax	[true   false]
NLS_LANG	Function	Specifies the language, territory, and character set of the client environment. The character set specified in NLS_LANG must match the character set of the terminal or terminal emulator. The character set specified in NLS_LANG can be different from the database character set, in which case Oracle automatically converts the character set.  See the <i>Oracle8i National Language Support Guide</i> for a list of values.
	Syntax	<i>language_territory.characterset</i>
	Example	<i>french_france.we8dec</i>
ORA_NLS33	Function	Specifies the directory where language, territory, character set, and linguistic definition files are stored.
	Syntax	<i>directory_path</i>
	Example	<i>\$ORACLE_HOME/ocommon/nls/admin/data</i>
ORACLE_BASE	Function	Specifies the base of the Oracle directory structure for Optimal Flexible Architecture (OFA) compliant databases.
	Syntax	<i>directory_path</i>
	Example	<i>/u01/app/oracle</i>
ORACLE_HOME	Function	Specifies the directory containing the Oracle software.
	Syntax	<i>directory_path</i>
	Example	<i>\$ORACLE_BASE/product/8.1.7</i>

**Table 1–1 Oracle8i Environment Variables on UNIX (Cont.)**

Variable	Detail	Definition
ORACLE_PATH	Function	Specifies the search path for files used by Oracle applications, such as *.sql (SQL*Plus), *.frm (Oracle Forms), and *.rpt (Oracle Reports). If the full path to the file is not specified, or is not in the current directory, the Oracle application uses ORACLE_PATH to locate the file.
	Syntax	Colon-separated list of directories <i>directory:directory:directory</i>
	Example	/u01/oracle/adhoc/8.1.7/bin: .  <b>Note:</b> The period adds the current working directory to the search path.
ORACLE_SID	Function	Specifies the Oracle system identifier.
	Syntax	A string of numbers and characters that must begin with a letter. Oracle Corporation recommends a maximum of eight characters. For more information, see the <i>Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX</i> .
	Example	SAL1
ORACLE_TRACE	Function	Turns on tracing of Bourne shell scripts during an installation. If set to T, many Oracle shell scripts run with the set -x flag on.
	Range of Values	T or not T.
ORAENV_ASK	Function	Controls whether coraenv or oraenv prompt for the value of the ORACLE_SID or ORACLE_HOME environment variables. If set to NO, they do not prompt; otherwise they do.
	Syntax	<i>string</i>
	Range of Values	NO or not NO.
SQLPATH	Function	Specifies the directory or list of directories that SQL*Plus searches for a login.sql file.
	Syntax	Colon-separated list of directories <i>directory:directory:directory</i>
	Example	/home:/home/oracle:/u01/oracle
TNS_ADMIN	Function	Specifies the directory containing the Net8 configuration files.
	Syntax	<i>directory_path</i>
	Range of Values	Any directory; for more information, see the <i>Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX</i> .
	Example	\$ORACLE_HOME/network/admin

**Table 1–1 Oracle8i Environment Variables on UNIX (Cont.)**

Variable	Detail	Definition
TWO_TASK	Function	Specifies the default Net8 connect string descriptor alias defined in the <code>tnsnames.ora</code> file.
	Syntax	Any available network alias.
	Range of Values	Any valid Net8 alias defined in the <code>tnsnames.ora</code> file.
	Example	PRODDB_TCP

## UNIX Environment Variables

[Table 1–2](#) provides the syntax for, and examples of, UNIX environment variables used with Oracle8i.

**Table 1–2 UNIX Environment Variables Used with Oracle8i**

Variable	Detail	Definition
CLASSPATH	Function	Used for Java applications. This variable differs with each Java application. See the product documentation for your Java application for more information.
	Syntax	<i>directory_path</i>
	Example	There is no default setting. CLASSPATH must include the <code>JRE_Location</code> , <code>\$ORACLE_HOME/product/jlib</code> directory where <code>JRE_Location</code> is defined as <code>\$ORACLE_HOME/JRE</code> .
DISPLAY	Function	Used by X-based tools. Specifies the display device used for input and output. See the X Windows documentation of the vendor for details.
	Syntax	<i>hostname:display</i> where the <i>hostname</i> is your system name (either IP address or alias); <i>display</i> is the monitor number. If you have a single monitor, the number is 0.
	Example	135.287.222.12:0 bambi:0
HOME	Function	The user's home directory.
LANG or LANGUAGE	Function	Specifies the language and character set used by the operating system for messages and other output. See the operating system documentation and the <i>Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX</i> .

**Table 1–2 UNIX Environment Variables Used with Oracle8i (Cont.)**

Variable	Detail	Definition
LPDEST	Function	Specifies the user's default printer.
	Syntax	<i>printer_name</i>
	Example	<i>doc_printer</i>
LD_LIBRARY64_PATH	Function	Specifies a search path for 64-bit shared libraries. Set LD_LIBRARY64_PATH to include the \$ORACLE_HOME/lib and \$ORACLE_HOME/javavm/admin directories. This environment variable is required if you are using Oracle products that use 64-bit shared libraries.
	Syntax	Colon-separated list of directories <i>directory:directory:directory</i>
	Example	<i>/usr/dt/lib:\$ORACLE_HOME/lib:\$ORACLE_HOME/javavm/admin</i>
LD_LIBRARYN32_PATH	Function	Specifies a search path for 32-bit shared libraries. Set LD_LIBRARYN32_PATH to include the \$ORACLE_HOME/lib32 directory. This environment variable is required if you are using Oracle products that use 32-bit shared libraries.
	Syntax	Colon-separated list of directories <i>directory:directory:directory</i>
	Example	<i>/usr/dt/lib:\$ORACLE_HOME/lib32</i>
LD_LIBRARY_PATH	Function	Used by the shared library loader ( <i>ld.so.1</i> ) at runtime to find shared object libraries. See the <i>ld.so.1</i> man pages for details.
	Syntax	Colon-separated list of directories <i>directory:directory:directory</i>
	Example	<i>/usr/dt/lib:\$ORACLE_HOME/lib</i>
PATH	Function	Used by the shell to locate executable programs; must include \$ORACLE_HOME/bin.
	Syntax	Colon-separated list of directories <i>directory:directory:directory</i>
	Example	<i>/bin:/usr/bin:/usr/local/bin: /usr/bin/X11:\$ORACLE_HOME/bin:\$HOME/bin.</i>
SGI_ABI	Function	<b>Note:</b> The period adds the current working directory to the search path. Ensures that the compiler and linker run in 64-bit mode whenever you compile or relink.
	Syntax	String (shell variable)
	Example	-64
PRINTER	Function	Selects the default printer.

**Table 1–2 UNIX Environment Variables Used with Oracle8i (Cont.)**

Variable	Detail	Definition
SHELL	Syntax	<i>printer_name</i>
	Example	doc_printer
	Function	Specifies the command interpreter used during a HOST command.
	Syntax	<i>shell_path</i>
	Range of Values	/bin/sh or /bin/csh or /bin/ksh or any other command interpreter supplied with SGI IRIX.
SRCHOME	Example	/bin/sh
	Function	Specifies the source directory for the installation. SRCHOME should be undefined when you run Oracle Universal Installer. If SRCHOME is defined, Oracle Universal Installer defaults to the location it specifies as the source of software to install.
	Syntax	<i>directory_path</i>
TERM	Example	\$ORACLE_HOME
	Function	Used by Oracle Toolkit II character mode tools and other UNIX tools to determine terminal types.
	Example	vt100
THREADS_FLAG	Function	Set the THREADS_FLAG environment variable to native. The default is <i>green</i> . SGI IRIX does not support mixing native and green threads.
	Syntax	String
	Example	\$THREADS_FLAG= native
TMPDIR	Function	Specifies the default directory for temporary disk files; if set, tools that create a temporary files do so in this directory.
	Syntax	<i>directory_path</i>
	Example	/u02/oracle/tmp
XENVIRONMENT	Function	Specifies a file containing X-Windows system resource definitions. See your X-Windows documentation for more information.

## Setting a Common Environment

The following section describes how to set a common UNIX environment.

### The oraenv Command File

The `oraenv` (`coraenv` for the C shell) command file is created during installation. It contains values for Oracle environment variables and provides:

- A central means of updating all user accounts with database changes
- A mechanism for switching between Oracle8i databases

You might find yourself frequently adding and removing databases from your development system or your users might be switching between several different Oracle databases installed on the same system. With `oraenv` each user shell startup file profile calls the `oraenv` command file.

---

---

**Note:** The C shell uses the `coraenv` command instead of the `oraenv` command.

---

---

### Local bin Directory

Place `oraenv` (or `coraenv`) and `dbhome` scripts in a local `bin` directory, separate from the Oracle software home directory, to ensure that these files are accessible to all users. Doing this also ensures that the `oraenv` script continues to work even if you change the path to specify a different Oracle home directory. The local `bin` directory is specified by the `root.sh` script, which you run after you install Oracle8i. The default location for the local `bin` directory on SGI IRIX is `/usr/local/bin`.

### Switching Between Databases

To switch from one database or database instance to another, call the `oraenv` routine. Reply to the prompt with the value of the `ORACLE_SID` environment variable of the database to which you are switching. Always provide the full path of the `oraenv` command file. For example:

```
$ . /usr/local/bin/oraenv
ORACLE_SID= [default]? sid
```

## Setting and Exporting the Value of a Variable in a Current Session

Use the `env` command to show the environment variable values that have been exported to the environment. The Bourne shell and Korn shell can set values without exporting them.

For the Bourne or Korn shell, enter:

```
$ ORACLE_SID=test
$ export ORACLE_SID
```

For the C shell, enter:

```
% setenv ORACLE_SID test
```

In the preceding example, *test* is the value of the `ORACLE_SID` environment variable.

## Setting the System Time

The `TZ` environment variable sets the time zone. It enables you to adjust the clock for daylight saving time changes or different time zones. The adjusted time is used to time-stamp files, produce the output of the `date` command, and obtain the current system date setting.

Oracle Corporation recommends that you do not change your personal `TZ` value. Using different values of `TZ` such as `GMT+24` might change the day a transaction is recorded. This affects Oracle applications that use `SYSDATE`, such as Oracle Financials. To avoid this problem, use sequence numbers to order a table instead of date columns.

## Product Executables

You can manually relink your product executables using a relink shell script located in the `$ORACLE_HOME/bin` directory. Relinking is necessary after applying any operating system patches or after an operating system upgrade.

The relink script manually relinks Oracle product executables depending on the products that have been installed in the Oracle home directory.

To relink product executables, enter the following command, where *parameter* is one of the values listed in [Table 1-3](#):

```
$ relink parameter
```

**Table 1-3 Relink Script Parameters**

Value	Description
all	Every product executable that has been installed
oracle	Oracle database executable only
network	net_client, net_server, nau, cman, cnames
client	net_client, otrace, plsql, client_sharedlib
interMedia	ctx, ordimg, ordaud, ordvir, md
precomp	All precompilers that have been installed
utilities	Utilities
oemagent	oemagent, odg

---

---

**Note:** Shut down Oracle Intelligent Agent and other Oracle programs in this Oracle home directory before relinking the database executable.

---

---

## System Global Area

The System Global Area (SGA) is the Oracle structure that is stored in shared memory. It contains static data structures, locks, and data buffers. Sufficient shared memory must be available to each Oracle process to address the entire SGA.

The maximum size of a single shared memory segment is specified by the SHMMAX parameter. For example, if SHMMAX is 132 MB and the SGA is 528 MB, the SGA requires four shared memory segments. This setting is located in the `/var/sysgen/stune` directory. The recommended value for SHMMAX is 4,294,967,296 regardless of the actual memory installed on the system.

If the size of the SGA exceeds the maximum size of a shared memory segment (SHMMAX), the value of the SHMSEG parameter is the maximum number of segments that can be attached by a process.

---

---

**Note:** Intimate Shared Memory (ISM) may cause problems when SHMMAX is smaller than the database SGA size.

---

---

Set the following initialization file parameters to control the size of the SGA:

- DB\_BLOCK\_BUFFERS
- DB\_BLOCK\_SIZE
- LOG\_BUFFER
- SORT\_AREA\_SIZE
- SHARED\_POOL\_SIZE
- JAVA\_POOL\_SIZE
- LARGE\_POOL\_SIZE

Use caution when setting values for these parameters. When values are set too high, too much of the computer's physical memory is devoted to shared memory, resulting in poor performance.

## Calculating the Size of the SGA

You can determine the SGA size in one of the following ways:

- Calculate the approximate size of an instance's SGA using the following formula:  
$$\begin{aligned} & (\text{DB\_BLOCK\_BUFFERS} \times \text{DB\_BLOCK\_SIZE}) \\ & + \text{SORT\_AREA\_SIZE} \\ & + \text{SHARED\_POOL\_SIZE} \\ & + \text{LOG\_BUFFER} \\ & + \text{JAVA\_POOL\_SIZE} \end{aligned}$$
- Display the size of the SGA for a running database using the SQL\*Plus SHOW SGA command. The result is shown in bytes.
- Determine the size of the SGA when you start your database system. The SGA size is displayed next to the heading Total System Global Area.

The address at which the SGA is attached affects the amount of virtual address space available for database buffers in the SGA and cursors in the user's application data area.

To relocate the SGA, perform the following steps:

1. Enter the following command:

```
$ tstshm
```

In the output from `tstshm`, the lines "Lowest shared memory address" and "Highest shared memory address" indicate the valid address range.

2. Check the "segment boundaries" output of the `tstshm` command to determine the valid virtual address boundaries at which a shared memory segment can be attached.
3. Change directory to the `$ORACLE_HOME/rdbms/lib` directory and enter the following command to generate the `ksms.s` file:

```
$ ORACLE_HOME/bin/genksms -b sgabeg > ksms.s
```

In this example, *sgabeg* is the starting address of the SGA (which defaults to `0x20000000`). It should fall within the range determined in step 2.

Never set the starting address below `0x01000000`. On most systems, this leaves approximately 7 MB for data segments. This amount must allow enough memory for the `SORT_AREA_SIZE` parameter and similar items.

With a start address of `0x1000000` you can achieve an overall SGA size of about 3.5 GB.

You might receive the following error messages if you reduced the value of the starting address:

- `ORA-4030: out of process memory when trying to allocate %s bytes (%s,%s)`
- `ORA-7324: smpall:malloc error while allocating sga.`

If you receive one of these messages, then you probably lowered the starting address into an area which the SGA needs. Increase the starting address, and try again.

4. Shut down the existing Oracle instance.
5. Enter the following command to rebuild the Oracle executable in the `$ORACLE_HOME/rdbms/lib` directory:

```
$ make -f ins_rdbms.mk ksms.o  
$ make -f ins_rdbms.mk ioracle
```

Using the `ioracle` target:

- Backs up the old executable (`oracle0`).
- Assigns the correct privileges to the new Oracle executable.
- Moves the new executable into the `$ORACLE_HOME/bin` directory.

The result is a new Oracle kernel that loads the SGA at the address specified when generating the `ksmc.s` file in step 3.

**See Also:** For more information on how the use of Java in the database affects SGA calculations, see the README file in the `$ORACLE_HOME/javavm/doc` directory.

## Oracle8i Memory Requirements and Usage

Calculate memory usage requirements to determine the number of users that the system can support. This also helps to determine the physical memory and swap space requirements.

1. To calculate the memory requirements, use the following formula:

$$\begin{aligned}
 & \text{size of the Oracle executable text} \\
 + & \text{size of the SGA} \\
 + & n * ( \text{size of tool executables private data section} \\
 & \quad + \text{size of the Oracle executables uninitialized data section} \\
 & \quad + 8192 \text{ bytes for the stack} \\
 & \quad + 2048 \text{ bytes for the processes user area} )
 \end{aligned}$$

In the preceding example, *n* represents the number of background processes.

2. For each client-server connection, use the following formula to estimate virtual memory requirements:

$$\begin{aligned}
 & \text{size of the Oracle executable data section} \\
 + & \text{size of the Oracle executables uninitialized data section} \\
 + & 8192 \text{ bytes for the stack} \\
 + & 2048 \text{ bytes for processes user area} \\
 + & \text{cursor area required for the application}
 \end{aligned}$$

3. Use the `size` command to estimate an executable's text size, private data section size, and uninitialized data section size (or DSS).

Program text is counted only once, no matter how many times the program is run, because the Oracle executable text is shared.

## Server Resource Limits

SGI IRIX inherits resource limits from the parent process (see `getrlimit(2)` in your operating system documentation). These limits apply to the Oracle8i shadow process that executes for user processes. The SGI IRIX default resource limits are high enough for any Oracle8i shadow or background process. However, if these limits are lowered, the Oracle8i system could be affected. Discuss this with your SGI IRIX system manager.

Disk quotas established for the `oracle` user can hinder the operation of the Oracle8i system. Confer with the SGI IRIX system manager before establishing disk quotas.

## Database Limits

[Table 1-4](#) lists the maximum and default values for parameters in a `CREATE DATABASE` or `CREATE CONTROLFILE` statement.

---

---

**Note:** Interdependencies between these parameters can affect allowable values.

---

---

**Table 1-4** *Determining the Size of Control Files*

Parameter	Default Value	Maximum Value
MAXDATAFILES	30	3000
MAXINSTANCES	1	63
MAXLOGFILES	16	255
MAXLOGHISTORY	100	65534
MAXLOGMEMBERS	2	5

**Table 1–5 Oracle-Specific File Size Limits**

File Type	Maximum Size
Datafiles where DB_BLOCK_SIZE is 2048	8,589,932,544
Datafiles where DB_BLOCK_SIZE is 4096	17,179,865,088
Datafiles where DB_BLOCK_SIZE is 8192	34,359,730,176
Datafiles where DB_BLOCK_SIZE is 16384	68,719,460,352
Datafiles where DB_BLOCK_SIZE is 32768	137,438,920,704

## Special Accounts and Groups

[Table 1–6](#) describes the special UNIX accounts required by the Oracle server.

**Table 1–6 Special UNIX Accounts**

Account	Description
oracle	The Oracle software owner represents the account that owns the Oracle8i software. This maintenance account requires DBA privileges to CREATE, STARTUP, SHUTDOWN, and CONNECT to the database as the INTERNAL user. The Oracle software owner must never be the root user.
root	The root user is a special UNIX account with maximum privileges (superuser). This account is used to configure the UNIX kernel, configure and install networking software, and create user accounts and groups.

[Table 1–7](#) describes the special Oracle Server accounts required by the Oracle server.

**Table 1–7 Oracle Server Accounts**

Account	Description
SYS	This is a standard Oracle8i account with DBA privileges automatically created during installation. The SYS account owns all the base tables for the data dictionary. This account is used by the DBA.
SYSTEM	This is a standard Oracle8i account with DBA privileges automatically created during installation. Additional tables or views can be created by the SYSTEM user. DBAs may log in as SYSTEM to monitor or maintain databases.

[Table 1–8](#) describes the special group accounts required by the Oracle server.

**Table 1–8 Special Group Accounts**

Group	Description
dba	The Oracle software owner is the only required member of the dba group. You can add any other UNIX user to the dba group. Members of this group have access to SQL*Plus specially privileged functions. If your account is not a member of the dba group, you must enter a password to connect as INTERNAL or gain access to the other administrative functions of SQL*Plus. The default OSDBA group is dba.
oinstall	All users installing Oracle8i in any Oracle home directory must belong to the same UNIX group. The Oracle Universal Installer inventory is shared by all Oracle home directories on a system and is group writable. Oracle recommends installing with oinstall as the primary group.
oper	This is an optional UNIX group. Members have database OPERATOR privileges. OPERATOR privileges are a restricted set of DBA privileges.
root	Only the root user should be a member of the root group.

## Security

Oracle8i uses features of the UNIX operating system to provide a secure environment for users. These features include file ownership, group accounts, and the ability of a program to change its user ID upon execution.

The two-task architecture of Oracle8i improves security by dividing work (and address space) between the user program and the Oracle program. All database access is achieved through the shadow process and special authorizations in the Oracle program.

**See Also:** For more information on security issues, see the *Oracle8i Administrator's Guide*.

---

## Groups and Security

Oracle programs are divided into two sets for security purposes: those executable by all users (*other*, in UNIX terms), and those executable by DBAs only. To improve security, Oracle Corporation recommends:

- The primary group for the `oracle` account should be the `oinstall` group.
- The `oracle` account must have the `dba` group as a secondary group.
- Only the `oracle` account should belong to the `oinstall` group.

## Security for Server Manager Commands

Oracle Corporation recommends that you restrict access to Server Manager. Only the Oracle software owner and `dba` group members should have access to the system privileges for `STARTUP`, `SHUTDOWN`, and `CONNECT AS INTERNAL`.

---

---

**Caution:** System-privileged statements can damage your database if used incorrectly. Note that users who are not in the `dba` group can connect as `INTERNAL` if they have the password.

---

---

## Security for Database Files

The user ID used to install Oracle8i owns the database files. The default user ID is `oracle`. Set the authorizations on these files to read/write by owner, and read-only for group or other users.

The Oracle software owner should own the directories containing the database files. For added security, revoke read permission from group and other users.

To access the protected database files, the Oracle program must have its set user ID, `setuid`, bit on.

The Oracle Universal Installer automatically sets the permissions of the Oracle executable to:

```
-rwsr-s--x 1 oracle dba
```

The `s` in the user execute field means that when you execute the Oracle program, it has an effective user ID of `oracle`, regardless of the actual user ID of the person starting it.

If you must manually set the permissions on the Oracle executable, enter:

```
$ chmod 6751 $ORACLE_HOME/bin/oracle
```

## Remote Passwords

You can administer a database from a remote computer such as a PC without having an operating system account. In this case, users are validated by using an Oracle8i password file, created and managed by the `orapwd` utility. You can also use password file validation on systems that support operating system accounts.

Local password files are located in the `$ORACLE_HOME/dbs` directory and contain the user name and password information for a single database. If there are multiple Oracle home directories on a system, each has a separate password file. To enable the database to use the password file, set the initialization file parameter `REMOTE_LOGIN_PASSWORDFILE` to `EXCLUSIVE`.

### Remote Authentication

The initialization file parameters shown in [Table 1-9](#) control the behavior of remote connections through non-secure protocols:

**Table 1-9 Parameters for Controlling Remote Connections**

Parameter	Description
<code>REMOTE_OS_AUTHENT</code>	Enables or disables <code>ops\$</code> connection
<code>OS_AUTHENT_PREFIX</code>	Used by <code>ops\$</code> accounts
<code>REMOTE_OS_ROLES</code>	Enables or disables roles through remote connections

**See Also:** For information on resource limits, see `getrlimit(2)` in your operating system documentation.

## Running the orapwd Utility

The `orapwd` utility, in the `$ORACLE_HOME/bin` directory, must be run by the Oracle software owner. The syntax of this command is:

```
$ orapwd file=filename password=password entries=max_users
```

Table 1–10 describes this syntax:

**Table 1–10 Syntax for Executing orapwd**

Variable	Description
<i>filename</i>	Name of the file where password information is written. The name of the file must be <code>orapwsid</code> , and you must supply the full pathname. Its contents are encrypted and not user-readable. This parameter is mandatory.
<i>password</i>	Initial password you selected for INTERNAL and SYS. Change this password after you create the database using an ALTER USER statement. This parameter is mandatory.
<i>max_users</i>	Maximum number of users allowed to connect to the database as SYSDBA or SYSOPER. This parameter is mandatory only if you want this password file to be EXCLUSIVE. Set <i>max_users</i> to a higher number than you expect to require because if you must exceed this value, you must create a new password file.

## Example of the orapwd Utility

The following is an example of the `orapwd` utility:

```
$ orapwd file=/u01/app/oracle/product/8.1.7/dbs/orapwV817
password=V817pw entries=30
```

## Customizing the `initsid.ora` File

The default `initsid.ora` file is provided with the Oracle8i software. The Oracle Universal Installer creates it in the `$ORACLE_BASE/admin/sid/pfile` directory. You can modify or customize the Oracle8i installation. A sample of the `initsid.ora` file is located in the `$ORACLE_HOME/dbs` directory.

Table 1–11 lists default initialization parameter values on SGI IRIX. All Oracle8i instances assume these values if you do not specify different values for them in the `initsid.ora` file. Oracle Corporation recommends that you include in the `initsid.ora` file only those parameters that differ from the default initialization parameter values.

Use the SQL\*Plus command SHOW PARAMETERS to display the current values of these parameters on the system.

**Table 1–11 Initialization Parameters**

Parameter	Default Value	Range of Values
BACKGROUND_DUMP_DEST	?/rdbsms/log	Valid directory name
BITMAP_MERGE_AREA_SIZE	1048576	65536 to unlimited
COMMIT_POINT_STRENGTH	1	0 to 255
CONTROL_FILES	?/dbs/cntrloracle_sid.dbf	Valid file names
CREATE_BITMAP_AREA_SIZE	8388608	65536 to unlimited
DB_BLOCK_BUFFERS	100 MB of buffers	4 MB to unlimited
DB_BLOCK_SIZE	2048	2 KB to 16 KB
DB_FILES	200	1 to 2000000
DB_FILE_DIRECT_IO_COUNT	64	0 to 1048576/block size
DB_FILE_MULTIBLOCK_READ_COUNT	8	1 to min(DB_BLOCK_BUFFERS/4, 1048576/DB_BLOCK_SIZE)
DISTRIBUTED_TRANSACTIONS	1/4 TRANSACTIONS	0 to unlimited
HASH_AREA_SIZE	2*SORT_AREA_SIZE	0 to unlimited
HASH_MULTIBLOCK_IO_COUNT	0 (self-tuned)	0 to min(127, DB_BLOCK_BUFFERS/4, 1048576/DB_BLOCK_SIZE)
JAVA_POOL_SIZE	20000000	1000000 to 1000000000
LOCK_SGA	FALSE	TRUE, FALSE
LOG_ARCHIVE_DEST	null	Valid directory names
LOG_ARCHIVE_FORMAT	"%t_%s.dbf"	Valid file names
LOG_BUFFER	512KB or 128KB multiplied by the number of systems, whichever is larger	66560 to unlimited
LOG_CHECKPOINT_INTERVAL	0	0 to unlimited
MTS_MAX_DISPATCHERS	5	The value of MTS_DISPATCHERS to the value of PROCESSES
MTS_MAX_SERVERS	2*MTS_SERVERS, if MTS_SERVERS > 20, else 20	The value of MTS_SERVERS to the value of PROCESSES
MTS_SERVERS	1, if MTS_DISPATCHERS is specified, else 0	1 to the value of PROCESSES
MTS_LISTENER_ADDRESS	ADDRESS=address	

**Table 1–11 Initialization Parameters (Cont.)**

Parameter	Default Value	Range of Values
NLS_LANGUAGE	AMERICAN	Valid language names
NLS_TERRITORY	AMERICA	Valid territory names
OBJECT_CACHE_MAX_SIZE_PERCENT	10	0 to unlimited
OBJECT_CACHE_OPTIMAL_SIZE	102400	10 KB to unlimited
OPEN_CURSORS	50	1 to unlimited
OS_AUTHENT_PREFIX	ops\$	Arbitrary string
PROCESSES	30, if not PARALLEL_AUTOMATIC_TUNING	6 to unlimited
SHARED_POOL_SIZE	64 MB on 64-bit systems, 8 MB on 32-bit systems	300000 to unlimited
SORT_AREA_SIZE	65536	0 to unlimited

**See Also:** For information on initialization parameters see *Oracle8i Server Reference*, *Oracle8i Administrator's Guide* and *Oracle8i Tuning*.

## Embedded PL/SQL Gateway

The embedded PL/SQL gateway is a gateway embedded in the Oracle8i server to provide native support for deploying database applications written in PL/SQL on the Web. The embedded PL/SQL gateway is implemented as an Oracle Servlet Engine (OSE) servlet, and relies upon the existence and configuration of both the OSE and mod\_ose, the Apache module which supports the OSE. The following instructions provide information on how to install and configure the gateway.

### Overview

Two Apache modules, mod\_ose and mod\_plsql, support web applications developed with PL/SQL.

The mod\_ose module acts as a request router for an OSE running within an Oracle8i instance. mod\_ose enables stateful OSE applications by routing stateful requests through the middle tier and back to a specified OSE/Oracle8i instance. Because the embedded PL/SQL gateway is implemented as an OSE servlet running in the Oracle8i server, it is able to host stateful, as well as stateless, PL/SQL web

applications. A stateful PL/SQL web application is one in which all database session states (for example, package and transaction) is preserved between requests.

The `mod_plsql` module is the name given to the PL/SQL gateway running within an Apache module in the middle tier server and executing PL/SQL procedures in a backend Oracle server using OCI. `mod_plsql` currently supports only stateless PL/SQL web applications.

**See Also:** For information on developing web applications with PL/SQL, refer to *Using mod\_plsql* which is generic PL/SQL gateway documentation.

## Installing the Embedded PL/SQL Gateway

As with all OSE servlets, the embedded PL/SQL gateway must be loaded and published. To load and publish the embedded PL/SQL gateway servlet:

1. Enter the following command to load the servlet, connect to SQL\*Plus as `sys`, and run the `initplgs.sql` SQL script:

```
SQL> @rdbms/admin/initplgs.sql
```

2. The name of the embedded PL/SQL gateway servlet is `oracle.plsql.web.PLSQLGatewayServlet`. To publish the servlet, enter the following command:

```
% $ORACLE_HOME/jis/bin/unix/sess_sh -s http://<OSE machine name>:<OSE port \
number -u sys/change_on_install -c "publishservlet -virtualpath \
pls/*/webdomains/contexts/default plsGateway \
SYS:oracle.plsql.web.PLSQLGatewayServlet"
```

This command publishes the gateway servlet as `plsGateway` with a default context. The servlet can be accessed using the virtual path `/pls`. The following example shows a URL that might access a gateway servlet:

```
http://dlsun240/pls/dadname/hello_world
```

**See Also:** For more information on using and publishing servlets, see the *Oracle Servlet Engine User's Guide*.

## Configuring Oracle PL/SQL Embedded Gateways

Configuration procedures for configuring the Apache server/mod\_ose and the embedded PL/SQL gateway are beyond the scope of this reference.

**See Also:** For information on configuring the Apache server/mod\_ose, see *Oracle Servlet Engine User's Guide*.

## Demonstrations Files

This section describes how to build and run the SQL\*Loader and PL/SQL demonstration programs installed with Oracle8i.

### SQL\*Loader Demonstrations

The following SQL\*Loader demonstration files are included with Oracle8i. Run the demonstrations in numerical order:

**Table 1–12 SQL\*Loader Demonstration Files**

ulcase1	ulcase3	ulcase5	ulcase7	ulcase9
ulcase2	ulcase4	ulcase6	ulcase8	

#### To create and run a demonstration

Run demonstrations while logged in as the user SCOTT/TIGER. Ensure that:

- The user SCOTT/TIGER has CONNECT and RESOURCE privileges.
- The EMP and DEPT tables exist and are empty.

In the following steps, *n* represents the demonstration number, listed in the previous section. To create and run a demonstration:

1. Run the `ulcasen.sql` script corresponding to the demonstration you want to run.

```
$ sqlplus SCOTT/TIGER @ulcasen.sql
```

2. Load the demonstration data into the objects:

```
$ sqlldr SCOTT/TIGER ulcasenctl
```

- For the `ulcase2` demonstration, you do not have to run the `ulcase2.sql` script.

- For the `ulcase6` demonstration, run the `ulcase6.sql` script, then enter the following at the command line:

```
$ sqlldr SCOTT/TIGER ulcase6 DIRECT=true
```

- For the `ulcase7` demonstration, run the `ulcase7s.sql` script, then enter the following at the command line:

```
$ sqlldr SCOTT/TIGER ulcase7
```

After running the demonstration, run `ulcase7e.sql` to drop the insert trigger and global variable package.

## Administering SQL\*Loader

Oracle8i incorporates SQL\*Loader functions. Demonstration and message files are in the `rdbms` directory.

The SQL\*Loader control file includes the following additional file processing option strings, the default being `str`, which takes no argument:

```
[ "str" | "fix n" | "var n" ]
```

String	Description
<code>str</code>	Specifies a stream of records, each terminated by a newline character, which are read in one record at a time. This is the default.
<code>fix</code>	Indicates that the file consists of fixed-length records, each of which is <i>n</i> bytes long, where <i>n</i> is an integer value.
<code>var</code>	Indicates that the file consists of variable-length records, each of which is <i>n</i> bytes long, where <i>n</i> is an integer value specified in the first five characters of the record.

If you do not select the file processing option, the information is processed by default as a stream of records (STR). You might find that `fix` mode yields faster performance than the default STR mode because it does not scan for record terminators.

## Newline Characters in Fixed Length Records

When using the FIX option to read a file containing fixed-length records, where each record is terminated by a newline character, include the length of the newline (one character) when specifying the record length to SQL \*Loader.

For example, to read the following file specify FIX 4 instead of FIX 3 to account for the additional newline character:

```
AAA newline
BBB newline
CCC newline
```

If you do not terminate the last record in a file of fixed records with a newline character, do not terminate the other records with a newline character either. Similarly, if you terminate the last record with a newline character, terminate all records with a newline character.

---

---

**Caution:** Certain text editors, such as `vi`, automatically terminate the last record of a file with a newline character. This leads to inconsistencies if the other records in the file are not terminated with newline characters.

---

---

## Removing Newline Characters

Use the `position(x:y)` function in the control file to discard the newline characters from fixed length records rather than loading them. For example, enter the following in your control file to discard newline characters from the fourth position:

```
load data
infile xyz.dat "fix 4"
into table abc
( dept position(01:03) char )
```

When this is done, newline characters are discarded because they are in the fourth position in each fixed-length record.

## PL/SQL Demonstrations

PL/SQL includes a number of sample programs that you can load. Demonstration and message files are in the `rdbms` directory. The Oracle8i database must be open and mounted to work with the sample programs:

1. Run SQL\*Plus and connect with the user/password SCOTT/TIGER:

```
$ cd $ORACLE_HOME/plsql/demo
$ sqlplus scott/tiger
```

2. To load the demonstrations, run `exampbld.sql` from SQL\*Plus:

```
SQL> @exampbld.sql
```

---



---

**Note:** Build the demonstrations in any Oracle account with sufficient permissions. Run the demonstrations in the same account as they were built.

---



---

The following kernel demonstrations are available:

```
examp1.sql      examp5.sql      examp11.sql     sample1.sql
examp2.sql      examp6.sql      examp12.sql     sample2.sql
examp3.sql      examp7.sql      examp13.sql     sample3.sql
examp4.sql      examp8.sql      examp14.sql     sample4.sql
extproc.sql
```

To run the PL/SQL demonstrations, run SQL\*Plus to connect to the database, using the same user/password used to create the demonstrations. Start the demonstration by typing an "at" sign (@) or the word START before the demonstration name. For example, to start the `examp1` demonstration, enter:

```
$ sqlplus SCOTT/TIGER
SQL> @examp1
```

To build the precompiler PL/SQL demonstrations, enter:

```
$ cd $ORACLE_HOME/plsql/demo
$ make -f demo_plsql.mk demos
```

[Table 1-13](#) lists the precompiler demonstrations.

**Table 1-13 Precompiler Demonstrations**

examp9.pc	examp10.pc	sample5.pc	sample6.pc
-----------	------------	------------	------------

If you want to build a single demonstration, enter its name as the argument in the `make` command. For example, to build the `examp9.pc` executable, enter:

```
$ make -f demo_plsql.mk examp9
```

To run the `examp9` demonstration from your current shell, enter:

```
$ ./examp9
```

To run the `extproc` demonstration, first add the following line to the file, `tnsnames.ora`:

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=plsf)))(CONNECT_DATA=(SID=extproc))
```

and the following line to the file, `listener.ora`:

```
SC=(SID_NAME=extproc)(ORACLE_HOME=/u01/app/oracle/product/8.1.7)
(PROGRAM=extproc)
```

then from your **SQL\*Plus** session, enter:

```
SQL> CONNECT SYSTEM/MANAGER
Connected.
SQL> GRANT CREATE LIBRARY TO SCOTT;
Grant succeeded.
SQL> CONNECT SCOTT/TIGER
Connected.
SQL> CREATE LIBRARY DEMOLIB AS
'$ORACLE_HOME/plsql/demo/extproc.so';
Library created.
```

Finally, to run the demonstrations:

```
SQL> CONNECT SCOTT/TIGER
Connected.
SQL> @extproc
```

## Database Examples

In the following examples, it is assumed that the local `bin` directory is `/usr/local/bin` and the production database is called `PROD`. In addition, `ORAENV_ASK` is reset to the default, `Yes`, after `oraenv` is executed. This ensures that the system prompts for a different `ORACLE_SID` environment variable the next time that the `oraenv` command is executed.

---

---

**Note:** Set the `ORAENV_ASK` environment variable to `no` to not prompt for the `ORACLE_SID` at startup.

---

---

If you have created a database manually instead of using Oracle Database Configuration Assistant, you must ensure that the system configuration is reflected in the `/var/opt/oracle/oratab` file.

For each server instance, add an entry in the following format:

```
ORACLE_SID:ORACLE_HOME:{Y|N}
```

The values `Y` or `N` indicate whether or not you want to activate the `dbstart` and `dbshut` scripts. The Oracle Database Configuration Assistant automatically adds an entry for each database it creates.

### Example of Single Instance

For the Bourne or Korn shell, add or replace the following line in the `.profile` file:

```
. local_bin_directory/oraenv
```

with the following lines:

```
PATH=${PATH}:/usr/local/bin
ORACLE_SID=PROD
export PATH ORACLE_SID
ORAENV_ASK=NO
export ORAENV_ASK
. oraenv
unset ORAENV_ASK=
```

For the C shell, add or replace the following line in the `.cshrc` file:

```
source local_bin_directory/coraenv
```

with the following lines:

```
setenv PATH ${PATH}:/usr/local/bin
setenv ORACLE_SID PROD
setenv ORAENV_ASK NO
source /usr/local/bin/coraenv
unset ORAENV_ASK
```

### Example of Multiple Instances

For multiple database instances, define the *sid* at startup.

For the Bourne or Korn shell, enter:

```
#!/usr/bin/sh
echo "The SIDs on this machine are:"
cat /var/opt/oracle/oratab | nawk-F: '{print $1}' | grep -v "#"
ORAENV_ASK="YES"
. /usr/local/bin/oraenv
```

For the C shell, enter:

```
#!/usr/bin/csh
echo "The SIDs on this machine are:"
cat /var/opt/oracle/oratab | nawk -F: '{print $1}' | grep -v "#"
set ORAENV_ASK="YES"
source /usr/local/bin/coraenv
```



---

---

# Tuning Oracle8i

The more your Oracle8i applications increase in complexity, the more you must tune the system to optimize performance and prevent data bottlenecks. This chapter describes how to configure your Oracle8i installation to optimize its performance. It contains the following sections:

- [The Importance of Tuning](#)
- [SGI IRIX Tools](#)
- [SQL Scripts](#)
- [Tuning Memory Management](#)
- [Tuning Disk I/O](#)
- [Monitoring Disk Performance](#)
- [Tuning CPU Usage](#)
- [Tuning Oracle Resource Contention and UNIX Kernel Parameters](#)
- [Tuning and Specifying Block Size and File Size](#)
- [Tuning the SGI IRIX Buffer Cache Size](#)
- [Tuning Oracle8i on IRIX Systems](#)
- [Using Trace and Alert Files](#)
- [Using Raw Devices/Volumes](#)

## The Importance of Tuning

Oracle8i is a highly-optimizable software product. Frequent tuning optimizes system performance and prevents data bottlenecks. Before tuning the system, observe its normal behavior using the SGI IRIX tools described in the next section.

**See Also:** For more information on tuning, see the *Oracle8i Designing and Tuning for Performance* guides.

## SGI IRIX Tools

SGI IRIX provides performance monitoring tools that you can use to assess database performance and determine database requirements. In addition to providing statistics for Oracle processes, these tools provide statistics for CPU usage, interrupts, swapping, paging, and context switching for the entire system.

**See Also:** For more information on UNIX tools, see *Oracle8i Designing and Tuning Performance*. SGI IRIX tools are described in the operating system documentation.

### sar

Use the `sar` command to monitor swapping, paging, disk, and CPU activity, depending on the switches that you supply with the command. The following command displays a summary of paging activity ten times, at ten second intervals:

```
$ sar -p 10 10
```

The following example shows sample output from the command:

```
14:14:55 atch/s pgin/s ppgin/s pflt/s vflt/s slock/s
14:15:05 0.00 0.00 0.00 0.60 1.00 0.00T
14:15:05 0.00 0.00 0.00 0.60 1.00 0.00T
14:15:15 0.00 0.00 0.00 0.10 0.60 0.00
14:15:25 0.00 0.00 0.00 0.00 0.00 0.00
14:15:35 0.00 0.00 0.00 0.00 0.00 0.00
14:15:45 0.00 0.00 0.00 0.00 0.00 0.00
14:15:55 0.00 0.00 0.00 0.00 0.00 0.00
14:16:05 0.00 0.00 0.00 0.00 0.00 0.00
14:16:15 0.00 0.00 0.00 0.00 0.00 0.00

Average 0.00 0.00 0.00 0.07 0.16 0.00
i
```

## swap

Use the following command to report information on swap space usage. A shortage of swap space can result in the system hanging and slow response time:

```
$ swap -l
      1 /dev/swap
          0,402    0    0 1048576 1026528 1048576    0
      2 /dev/dsk/dks0d1s3
          0,408    0    0 7168000 7147520 7168000    0
      3 /dev/dsk/dks0d1s4
          0,411    0    0 7168000 7151488 7168000    0
      4 /dev/dsk/dks0d1s5
          0,414    0    0 2662400 2642240 2662400    0
```

## SQL Scripts

Oracle8i release 3 (8.1.7) includes a set of packages for database tuning called STATPACKS. For more information on STATPACKS, see *Oracle8i Designing and Tuning Performance*.

The `utlbstat.sql` and `utlestat.sql` scripts are used to monitor Oracle database performance and tune the System Global Area (SGA) data structures. For more information on these scripts, see *Oracle8i Designing and Tuning for Performance*. On SGI IRIX, the scripts are located in the `$ORACLE_HOME/rdbms/admin/` directory.

## Tuning Memory Management

Start the memory tuning process by measuring paging and swapping space to determine how much memory is available.

The Oracle buffer manager ensures that the more frequently accessed data is cached longer. Monitoring the buffer manager and tuning the buffer cache can have a significant influence on Oracle performance. The optimal Oracle buffer size for your system depends on the overall system load and the relative priority of Oracle over other applications.

## Allocate Sufficient Swap Space

Try to minimize swapping because it causes significant UNIX overhead. Use the `sar -w` command on SGI IRIX to check for swapping.

If your system is swapping and you must conserve memory:

- Avoid running unnecessary system daemon processes or application processes
- Decrease the number of database buffers to free some memory
- Decrease the number of UNIX file buffers, especially if you are using raw devices

On SGI IRIX use the `swap -l` command to determine how much swap space is currently in use. Use the `swap -a` command to add swap space to your system. Consult your SGI IRIX documentation for further information.

## Control Paging

Paging might not present as serious a problem as swapping, because an entire program does not have to be stored in memory to run. A small number of page-outs might not noticeably affect the performance of your system.

To detect excessive paging, run measurements during periods of fast response or idle time to compare against measurements from periods of slow response.

Use the `sar -p` command to monitor paging. The following columns from the output of this command are important:

<code>vflt/s</code>	Indicates the number of address translation page faults. Address translation faults occur when a process references a valid page not in memory.
<code>rclm/s</code>	Indicates the number of valid pages that have been reclaimed and added to the free list by page-out activity. This value should be zero.

If your system consistently has excessive page-out activity, consider the following solutions:

- Install more memory
- Move some of the work to another system
- Configure your kernel to use less memory

## Hold the SGA in a Single Shared Memory Segment

You cannot start the database without sufficient shared memory. If necessary, reconfigure the UNIX kernel to increase shared memory.

**See Also:** For more information on the SGA, see "[System Global Area](#)" on page 1-10 and the *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX*.

## Tuning Disk I/O

I/O bottlenecks are the easiest performance problems to identify. Balance I/O evenly across all available disks to reduce disk access times. For smaller databases and those not using the Parallel Query option, ensure that different datafiles and tablespaces are distributed across the available disks.

## Tune the Database Writer to Increase Write Bandwidth

Oracle offers asynchronous I/O, multiple DBWR processes, and I/O slaves as solutions to prevent database writer (DBWR) activity from becoming a bottleneck.

### Asynchronous I/O

Asynchronous I/O allows processes to proceed with the next operation without having to wait after issuing a write and therefore improves system performance by minimizing idle time. SGI IRIX supports asynchronous I/O to both raw partitions, XLV or XVM volumes, and filesystem datafiles.

### I/O Slaves

I/O slaves are specialized processes whose only function is to perform I/O. They replace the Oracle7 feature, Multiple DBWRs. In fact, they are a generalization of Multiple DBWRs and can be deployed by other processes as well. They can operate whether or not asynchronous I/O is available. They are allocated memory from the LARGE\_POOL\_SIZE parameters, if set, otherwise they are allocated memory from shared memory buffers. I/O slaves include a set of initialization parameters that allow a degree of control over the way they operate.

[Table 2-1](#) lists the initialization parameters that control the operation of asynchronous I/O and I/O slaves.

**Table 2-1 Initialization Parameters for Asynchronous I/O and I/O Slaves**

Parameter	Range of Values	Default Value
DISK_ASYNCH_IO	TRUE/FALSE	TRUE
TAPE_ASYNCH_IO	TRUE/FALSE	TRUE
BACKUP_DISK_IO_SLAVES	TRUE/FALSE	FALSE
BACKUP_TAPE_IO_SLAVES	TRUE/FALSE	FALSE
DBWR_IO_SLAVES	0 - 999	0
ARCH_IO_SLAVES	0 - 999	0
DB_WRITER_PROCESSES	1-10	1

There might be times when the use of asynchronous I/O is not desirable or not possible. The first two parameters in [Table 2-1](#), DISK\_ASYNCH\_IO and TAPE\_ASYNCH\_IO, allow asynchronous I/O to be switched off respectively for disk and tape devices. Because the number of I/O slaves for each process type defaults to zero, no I/O slaves are deployed unless set.

Set the DBWR\_IO\_SLAVES parameter to greater than 0 if DISK\_ASYNCH\_IO or TAPE\_ASYNCH\_IO is disabled, otherwise DBWR becomes a bottleneck. In this case, the optimal value on SGI IRIX for DBWR\_IO\_SLAVES is 4.

DB\_WRITER\_PROCESSES replaces the Oracle7 parameter DB\_WRITERS and specifies the initial number of database writer processes for an instance. If you use DBWR\_IO\_SLAVES, only one database writer process is used, regardless of the setting for DB\_WRITER\_PROCESSES.

See "[Customizing the initsid.ora File](#)" on page 1-19 for information on other initialization parameters.

## Choose the Appropriate File System Type

SGI IRIX allows a choice of file systems. File systems have different characteristics, and the techniques they use to access data can have a substantial affect on database performance. Typical file system choices are:

- The EFS file system
- The XFS file system
- Data access through an XLV device
- Data access through raw device drivers

## Monitoring Disk Performance

To monitor disk performance, use the `sar -d` command. [Table 2-2](#) describes important `sar -d` output fields.

**Table 2-2** *Important sar -d Output Fields*

Field	Description
%busy	Percentage of time device was busy performing I/O requests
avque	Average number of I/O requests outstanding
r+w/s	Number of read and write I/O requests for each second
await	Average time in milliseconds that I/O requests wait idly on queue
avserv	Average time in milliseconds to service I/O requests (which for disks includes seek, rotational latency, and data transfer time)

High values for %busy, avque, await, and avserv can indicate I/O bottlenecks. Correlate busy devices from the `sar -d` output with the Oracle datafiles stored on those devices. Determine which Oracle datafiles are causing the I/O bottleneck and spread these datafiles across additional devices.

## Disk Performance Issues

Oracle block sizes should either match disk block sizes or be a multiple of disk block size.

If possible, do a file system check on the partition before using it for database files, then make a new file system to ensure that it is clean and unfragmented. Distribute disk I/O as evenly as possible, and separate log files from database files.

## Tuning CPU Usage

The following sections describe how to tune CPU usage.

### Keep All Oracle Users and Processes at the Same Priority

Oracle is designed to operate with all users and background processes operating at the same priority level. Changing priorities cause unexpected effects on contention and response times.

For example, if the log writer process (LGWR) gets a low priority, it is not executed frequently enough and LGWR becomes a bottleneck. On the other hand, if LGWR has a high priority, user processes can suffer poor response time.

### Use Processor Affinity and Binding on Multi-Processor Systems

In a multi-processor environment, use processor affinity and binding if it is available on your system. Processor binding prevents a process from migrating from one CPU to another, allowing the information in the CPU cache to be better utilized. You can bind a server shadow process to make use of the cache as it is always active, and let background processes flow between CPUs.

### Use a Client/Server Configuration

If your system is CPU-bound, move applications to a separate system to reduce the load on the CPU. For example, you can off-load foreground processes such as Oracle Forms to a client system to free CPU cycles on the database server system.

### Use the Post-Wait Driver

Oracle processes usually use semaphores to coordinate access to shared resources. If a shared resource is locked, a process suspends and waits for the shared resource to become available.

One way to improve shared resource coordination is to use a post-wait driver instead of semaphores. A post-wait driver is a faster, less expensive synchronization mechanism than a semaphore. The Oracle Post-Wait driver implements an optimized mechanism of inter-process communication, without the overhead of signal handlers or semaphores. It improves performance for Oracle8i.

To enable the Oracle Post-Wait driver, set the `ORA_USE_PW` environment variable to 1 before starting the database:

```
$ ORA_USE_PW= 1
```

---

---

**Note:** If you intend to run more than four Oracle databases on your system, or if you intend to have more than 512 back-end server processes for a database, you must modify the `/var/sysgen/master.d/postwait` system file. If you modify this file, you must re-generate the UNIX kernel. See the SGI IRIX operating system documentation for information on re-generating the UNIX kernel.

---

---

## Use Single-Task Linking for Large Exports/Imports and SQL\*Loader Jobs

If you must transfer large amounts of data between the user and Oracle8i (for example, using the `export` and `import` utilities), it is efficient to use single-task architecture. To make the single-task import (`impst`), export (`expst`), and SQL\*Loader (`sqlldrst`) executables, use the `ins_rdbms.mk` makefile in the `$ORACLE_HOME/rdbms/lib` directory. The following example makes the `impst`, `expst`, and `sqlldrst` executables:

```
$ cd $ORACLE_HOME/rdbms/lib
$ make -f ins_rdbms.mk singletask
```

---

---

**Note:** Linking Oracle executables as a single-task allows a user process to directly access the entire SGA. Running single-task executables requires more memory because the Oracle executable text is no longer shared between the front-end and background processes.

---

---

## Tuning Oracle Resource Contention and UNIX Kernel Parameters

You can improve performance by keeping the UNIX kernel as small as possible. The UNIX kernel typically pre-allocates physical RAM, leaving less memory available for other processes such as Oracle.

Traditionally, kernel parameters such as `NBUF`, `NFILE`, and `NOFILES` were used to adjust kernel size. However, most UNIX implementations dynamically adjust those

parameters at run time, even though they are present in the UNIX configuration file.

Look for memory-mapped video drivers, networking drivers, and disk drivers that could be de-installed, freeing more memory for use by other processes.

---

---

**Note:** Remember to make a backup copy of your UNIX kernel. See your operating system documentation for additional information.

---

---

## Tuning and Specifying Block Size and File Size

This section describes how you can improve the performance of Oracle8i by optimizing the size of Oracle blocks for the files in your database.

---

---

**Note:** To change block size, you must create a new database. To determine the most efficient configuration, experiment with block size before transferring your data to the new database.

---

---

On SGI IRIX, the default Oracle block size is 2 KB and the maximum block size is 32 KB. You can set the block size to any multiple of 2 KB up to a maximum of 16 KB, inclusive.

The optimal block size is typically the default size, however, it depends on the applications. To create a database with a different Oracle block size, add the following line to the `initsid.ora` file before creating the database:

```
db_block_size=new_block_size
```

---

---

**Note:** The value that you choose for the `DB_BLOCK_SIZE` parameter determines the maximum size of certain types of Oracle files. See [Table 1-5](#) for more information on file size limits.

---

---

## Tuning the SGI IRIX Buffer Cache Size

To take full advantage of raw devices, adjust the size of the Oracle8i buffer cache and, if memory is limited, the SGI IRIX buffer cache:

The SGI IRIX buffer cache holds blocks of data in memory while they are being transferred from memory to disk, or vice versa.

The Oracle8i buffer cache is the area in memory that stores the Oracle database buffers. Since Oracle8i can use raw devices, it does not need to use the SGI IRIX buffer cache.

If you use raw devices, you must increase the size of the Oracle8i buffer cache. If the amount of memory on the system is limited, make a corresponding decrease in the SGI IRIX buffer cache size. It is possible to increase or decrease the Oracle8i Buffer Cache by modifying the `DB_BLOCK_BUFFERS` parameter in the `initsid.ora` file and restarting the instance.

Use the `sar` command to determine which buffer caches you must increase or decrease. [Table 2-3](#) shows the options of the `sar` command.

**Table 2-3** *sar* Command options

Option	Description
-b	Reports the SGI IRIX buffer cache activity
-w	Reports the SGI IRIX swapping activity
-u	Reports CPU utilization
-r	Reports memory utilization
-p	Reports the SGI IRIX paging activity

### Adjusting Cache Size

To adjust the cache size, perform one of the following:

- Increase the Oracle8i buffer cache size while the cache hit ratio continues to increase.
- Decrease cache sizes if the swapping/paging activity becomes high.

## Tuning Oracle8i on IRIX Systems

The following sections describe the `IRIX_CPU_AFFNITY` and `IRIX_SCHEDULER` initialization parameters.

## IRIX\_CPU\_AFFINITY

The `IRIX_CPU_AFFINITY` initialization parameter enables you to specify processor affinity on an IRIX system. This parameter provides more control over system resources. The syntax is as follows:

```
IRIX_CPU_AFFINITY="NAME=MASK [NAME=MASK...]"
```

In the preceding example:

**MASK** is a single CPU number or a range of CPU numbers or the special keyword 'ANY' which denotes the `sysmp(2)` option `MP_RUNANYWHERE`. Specify a range of CPU numbers in the form *start\_number:end\_number* where the *start\_number* is less than the *end\_number*.

**NAME** is one or more of the following values:

Value	Description
ALL	Apply the MASK to all Oracle8i processes not included by a previously specified NAME=MASK combination.
OTHER	Apply the MASK to all Oracle8i processes not covered by the sys NAME.
SYS	Apply the MASK to any Oracle8i system processes not covered by another specified NAME=MASK combination.
PMON	Apply the MASK to the Process Monitor process.
DBWR	Apply the MASK to the Database Writer process.
SMON	Apply the MASK to the System Monitor process.
CKPT	Apply the MASK to the Checkpoint process.
ARCH	Apply the MASK to the Archiver process.
RECO	Apply the MASK to the Recovered process.
LOCK	Apply the MASK to the Lock Manager process.
JOBQ	Apply the MASK to the Job Queue process.
AQTM	Apply the MASK to the Advanced Queueing Time Monitor process.
IOS	Apply the MASK to the IO Slave processes.

---



---

**Note:** The default value is `ALL=ANY`

---



---

### Examples

The following specification of the `IRIX_CPU_AFFINITY` parameter assigns the PMON, SMON, and CKPT processes to CPU 1, the DBWR process to CPU 2, and the LOCK process is to CPU 3:

```
IRIX_CPU_AFFINITY="pmon=1 dbwr=2 smon=1 lock=3 ckpt=1 other=4:7"
```

All other Oracle8i processes are assigned in round-robin fashion (based on Oracle process id) to CPU 4 through 7.

The following specification of the `IRIX_CPU_AFFINITY` parameter assigns all the Oracle8i system processes to CPUs 0 through 4 in round-robin fashion based on the Oracle process id and assigns all other Oracle8i processes to CPUs 5 through 19 in round-robin fashion based on the Oracle process id:

```
IRIX_CPU_AFFINITY="sys=0:4 other=5:19"
```

### IRIX\_SCHEDULER

The `IRIX_SCHEDULER` initialization parameter allows you to specify certain IRIX scheduler options.

Default value is none.

The syntax is as follows:

```
IRIX_SCHEDULER="NAME=MASK [NAME=MASK...]"
```

In the preceding example:

**MASK** is a numeric value valid for the particular **NAME** value

**NAME** is one or more of the following values

Value	Description
<code>renice</code>	Apply the MASK as the <code>schedctl(2)</code> RENICE value.
<code>slice</code>	Apply the MASK as the <code>schedctl(2)</code> SLICE value.

The `$ORACLE_HOME/bin/oracle` executable might require you to set super user ID `root` privilege, depending on the scheduler options used.

See the `schedctl` man page for the details of the `schedctl` command for information on the `LOCK_SGA` parameter.

### Examples

The following examples demonstrate using the `IRIX_SCHEDULER` parameter:

```
IRIX_SCHEDULER="slice=50"
```

```
IRIX_SCHEDULER="renice=39 slice=50"
```

## Lock the SGA in Physical Memory

The primary function of the SGA is to cache database information. If the SGA begins paging to disk, caching becomes an overhead rather than a benefit. If you have installed Oracle8i on SGI IRIX, you can use the `LOCK_SGA` parameter to lock the memory pages associated with the SGA. Such locking prevents paging and helps asynchronous I/O on raw files to work efficiently. Locked memory is not available for use by other applications.

This option should be used only if there is enough physical memory on the system to support the Oracle instance, the applications, and other users. If the amount of physical memory on the system is insufficient, performance degradation may occur due to increased memory paging or swapping.

To lock the SGA:

1. Use a text editor to open the `initsid.ora` file.
2. Enter the following on a blank line in the `initsid.ora` file:  

```
LOCK_SGA=TRUE
```
3. Save and close the file.
4. Log in as `root`.
5. Make sure that the `MAXLKMEM` system parameter is set to at least 2050. Type the following at the command prompt:

```
$ systune maxlkmem 2050
```

The following message appears:

```
maxlkmem=current_setting  
Do you really want to change maxlkmem to 2050?
```

6. If the current setting is 2050, type N to preserve the setting, otherwise type Y to change the setting to 2050.

Any changes made in this way do not affect the kernel default setting for this parameter. For more information, type `man systune(1M)` at the command prompt or consult your system administrator.

## Using Trace and Alert Files

This section describes the trace (or dump) and alert files Oracle8i creates to diagnose and resolve operating problems.

### Trace File Names

The file name format of a trace file is *processname\_sid\_unixpid.trc*, where:

<i>processname</i>	Is a three- or four-character process name showing which Oracle8i process the trace file is from (for example, <i>pmon</i> , <i>dbwr</i> , <i>ora</i> , or <i>reco</i> )
<i>sid</i>	Is the instance system identifier
<i>unixpid</i>	Is the UNIX process ID number
<i>.trc</i>	Is a filename extension appended to all trace file names

### Alert Files

The *alert\_sid.log* file is associated with a database and is located in the directory specified by the *initsid.ora* parameter `BACKGROUND_DUMP_DEST`. The default directory is `$ORACLE_HOME/rdbms/log`.

## Using Raw Devices/Volumes

This section describes the use of raw devices on Oracle8i.

### Disadvantages of Raw Devices/Volumes

Raw devices/volumes have the following disadvantages when used on SGI IRIX:

- They may not solve problems with ULIMIT that can arise when exporting tables larger than 1 MB (such as another disk partition).
- When raw devices and operating system files are mixed within an Oracle8i database, the operating system files must still be within the value of the ULIMIT parameter.
- They may not solve problems with ULIMIT that can arise when reading the contents of the Oracle distribution media onto the disk.
- Small client systems usually cannot use sufficiently large raw device/volume partitions. Disk partitions usually come in odd sizes that do not lend themselves to good database architecture.
- If a particular disk drive has intense I/O activity and performance would benefit from movement of an Oracle data file to another drive, it is likely that no acceptably sized section exists on a drive with less I/O activity. Moving data files around, a common advantage of UNIX, may not be possible with raw devices/volumes.
- Adding space to a tablespace can be a difficult process in a raw device/volume environment. Occasionally, all raw partitions are assigned data files at initial configuration time, leaving no raw storage to accommodate normal tablespace growth.

## Guidelines for Using Raw Devices/Volumes

In addition to the disadvantages described in the previous section, you should consider the following issues when deciding whether to use raw devices/volumes:

- [Raw Disk Partition Availability](#)
- [Logical Volume Manager](#)
- [Dynamic Performance Tuning](#)
- [Mirroring and Online Disk Replacement](#)

### Raw Disk Partition Availability

Use raw devices/volumes for Oracle files only if your site has at least as many raw disk partitions as Oracle datafiles. If the raw disk partitions are already formatted, match datafile size to partition size as closely as possible to avoid wasting space.

### Logical Volume Manager

With logical volumes, you can create logical disks based on raw partition availability. Because logical disks can be moved to more than one disk, the disk drives do not have to be reformatted to obtain logical disk sizes.

**See Also:** To review the latest features of the SGI Logical Volume Manager, see the *XVM Volume Manager Administrator's Guide* and the *IRIX Admin: Disks and Filesystems* guide.

### Dynamic Performance Tuning

You can optimize disk performance when the database is online by moving hot spots to cooler drives.

Most hardware vendors who provide the logical disk facility also provide a graphical user interface that can be used for tuning.

### Mirroring and Online Disk Replacement

You can mirror logical volumes to protect against loss of data. If one copy of a mirror fails, dynamic re-synchronization is possible.

Some vendors also provide the ability to replace drives online in conjunction with the mirroring facility.

## Raw Device Setup

Consider the following items when creating raw devices:

- Ensure that the owner and group are `oracle` and `oinstall`, respectively. .
- The size of an Oracle datafile created in a raw partition must be at least two Oracle block sizes smaller than the size of the raw partition.

---

---

**Caution:** Do not attempt to set up raw devices without the help of an experienced system administrator and specific knowledge about the system you are using.

---

---

---

---

# Administering SQL\*Plus

This chapter describes how to use and administer SQL\*Plus on Oracle8i. It contains the following sections:

- [Administering SQL\\*Plus](#)
- [Using SQL\\*Plus](#)
- [Restrictions of SQL\\*Plus](#)

## Administering SQL\*Plus

This section describes how to administer SQL\*Plus.

### Using Setup Files

When you start SQL\*Plus, it executes the `glogin.sql` site profile setup file and then executes the `login.sql` user profile setup file.

#### Using the Site Profile File

The global site profile file is the `$ORACLE_HOME/sqlplus/admin/glogin.sql` directory. The default site profile is placed in the `$ORACLE_HOME/sqlplus/admin` directory when SQL\*Plus is installed. If a site profile already exists, it is overwritten. If SQL\*Plus is de-installed, the site profile file is deleted.

#### Using the User Profile File

The user profile file is `login.sql`. SQL\*Plus looks in the current directory, and then in the directories you specify until it finds the `login.sql` file. You can specify the directories to search by setting the `SQLPATH` environment variable to a colon-separated list of directories.

For example, if the current directory is `/u02/oracle` and `SQLPATH` is set to `/home:/home/oracle:/u01/oracle`, SQL\*Plus looks for the `login.sql` file in the following order:

1. `/u02/oracle` (the current directory)
2. `/home`
3. `/home/oracle`
4. `/u01/oracle`

The options set in the `login.sql` file override those set in the `glogin.sql` file.

**See Also:** For more information on profile files, see the *SQL\*Plus User's Guide and Reference*.

## Using the PRODUCT\_USER\_PROFILE Table

During a Typical installation, the PRODUCT\_USER\_PROFILE table is created automatically. This table is used to disable the SQL and SQL\*Plus commands you specify. To recreate this table, run the \$ORACLE\_HOME/sqlplus/admin/pupbld.sql script in the SYSTEM schema.

For example, enter:

```
$ sqlplus SYSTEM/MANAGER
SQL> @?/sqlplus/admin/pupbld.sql
```

SQL\*Plus uses the value of the ORACLE\_HOME environment variable wherever a question mark (?) appears.

## Using Demonstration Tables

SQL\*Plus is shipped with demonstration tables that you can use for testing.

### Performing a Typical Installation

During a Typical installation, the user SCOTT and the demonstration tables are created automatically.

### Creating Demonstration Tables Manually

Use the \$ORACLE\_HOME/sqlplus/demo/demobld.sql SQL script to create the demonstration tables. In SQL\*Plus, you can use any user name to run the demobld.sql file to create the demonstration tables in a schema. For example, enter:

```
$ sqlplus SCOTT/TIGER
SQL> @?/sqlplus/demo/demobld.sql
```

You can also use the \$ORACLE\_HOME/bin/demobld shell script to run the \$ORACLE\_HOME/sqlplus/demo/demobld.sql script by entering:

```
$ demobld scott tiger
```

## Deleting Demonstration Tables

Use the `$ORACLE_HOME/sqlplus/demo/demodrop.sql` script to drop demonstration tables. In SQL\*Plus, you can use any user name to drop the demonstration tables in the user's schema. For example, enter:

```
$ SQLPLUS SCOTT/TIGER
SQL> @?/sqlplus/demo/demodrop.sql
```

You can also use the `$ORACLE_HOME/bin/demodrop` shell script to run the `$ORACLE_HOME/sqlplus/demo/demodrop.sql` script by entering:

```
$ demodrop scott tiger
```

---

---

**Note:** Both the `demobld.sql` and `demodrop.sql` scripts drop the EMP, DEPT, BONUS, SALGRADE, and DUMMY tables. Before to running either script, make sure that no table with one of these names exists in the desired schema, or the table data is lost.

---

---

## Installing and Removing the Help Facility

This section describes how to install and remove the help facility.

### Performing a Typical Installation

When you copy a starter database with pre-built datafiles as part of the Typical installation or as an option in Oracle Database Configuration Assistant, SQL\*Plus automatically installs the Help Facility.

### Using the Database Configuration Assistant

You can use the Oracle Database Configuration Assistant to create help tables when creating a database.

### Installing the Help Facility Manually

You can use the `$ORACLE_HOME/bin/helpins` shell script to manually install the help facility. Before you run the script, set the `SYSTEM_PASS` environment variable to the `SYSTEM` schema name and password. For example, enter:

```
$ setenv SYSTEM_PASS SYSTEM/MANAGER
$ helpins
```

---

If the `SYSTEM_PASS` variable is not set, the `helpins` scripts prompt you for the `SYSTEM` password and load the help data into the `SYSTEM` schema. Run the `$ORACLE_HOME/sqlplus/help/helpbld.sql` script with the `helpus.sql` script to create the help facility tables. Enter the following:

```
$ sqlplus SYSTEM/MANAGER
SQL> @?/sqlplus/admin/help/helpbld.sql
```

---

---

**Note:** Both the `helpins` shell script and the `helpbld.sql` `SQL*Plus` script drop existing help facility tables before creating new tables.

---

---

For example, enter:

```
$ sqlplus SYSTEM/MANAGER
SQL> @?/sqlplus/admin/help/helpbld.sql
```

---

---

**Note:** Both the `helpins` shell script and the `helpbld.sql` `SQL*Plus` script drop existing help facility tables before creating new tables.

---

---

You can also run `$ORACLE_HOME/sqlplus/help/helpdrop.sql` in `SQL*Plus` to manually drop the help facility tables in a schema. For example, enter:

```
$ sqlplus SYSTEM/MANAGER
SQL> @?/sqlplus/admin/help/helpdrop.sql
```

**See Also:** For more information on the help facility, see the *SQL\*Plus User's Guide and Reference*.

## Using SQL\*Plus

This section describes how to use `SQL*Plus`.

### Using a System Editor from SQL\*Plus

If you enter an `ED` or `EDIT` command at the `SQL*Plus` prompt, the system starts an operating system editor, such as `ed`, `emacs`, `ned`, or `vi`. Your `PATH` variable must include the directory of the editor.

When you start the editor, the current SQL buffer is placed in the editor. When you exit the editor, the changed SQL buffer is returned to SQL\*Plus.

You can specify which editor starts by defining the SQL\*Plus `_editor` variable. This variable can be set in `glogin.sql`, in `login.sql`, or entered during a SQL\*Plus session. For example, to set the default editor to `vi`, enter:

```
SQL> DEFINE_EDITOR=VI
```

If you do not set the `_editor` variable, the value of either the `EDITOR` or `VISUAL` environment variable is used. If both are set, the `EDITOR` variable value is used. When `_editor`, `EDITOR`, and `VISUAL` are not specified, the default editor is `ed`.

If you run the editor, SQL\*Plus uses the `afiedt.buf` temporary file to pass text to the editor. You can use the `SET EDITFILE` command to rename this file. For example, enter:

```
SQL> SET EDITFILE /tmp/myfile.sql
```

SQL\*Plus does not delete the temporary file.

## Running Operating System Commands from SQL\*Plus

Using the `HOST` command or an exclamation point (!) as the first character after the SQL\*Plus prompt causes subsequent characters to be passed to a sub-shell. The `SHELL` environment variable sets the shell used to execute operating system commands. The default shell is `/bin/sh`. If the shell cannot be executed, an error message is displayed.

To return to SQL\*Plus, enter `exit` or press `Ctrl+D`.

For example, to enter one command, enter:

```
SQL>! COMMAND
```

In the preceding example, `command` represents the operating system command you want to execute.

To enter multiple operating system commands from SQL\*Plus, enter the `HOST` or ! command then press return. SQL\*Plus returns you to the operating system prompt.

## Interrupting SQL\*Plus

While running SQL\*Plus, you can stop the scrolling record display and terminate a SQL statement by pressing `[Ctrl]+[c]`.

## Using the SPOOL Command

The default extension name of files generated by the SPOOL command is `.lst`. To change this extension, specify a spool file containing a period (`.`). For example, enter:

```
SQL> SPOOL query.txt
```

## Restrictions of SQL\*Plus

This section describes SQL\*Plus restrictions.

## Resizing Windows

The default values for SQL\*Plus `LINESIZE` and `PAGESIZE` do not automatically adjust for window size.

## Return Codes

UNIX return codes use only one byte, which is not enough space to return an Oracle error code. The range for a return code is 0 to 255.



---

---

# Using Oracle Precompilers and the Oracle Call Interface

This chapter describes Oracle Precompilers and the Oracle Call Interface. It contains the following sections:

- [Overview of Oracle Precompilers](#)
- [Pro\\*C/C++ Precompiler](#)
- [Pro\\*FORTRAN Precompiler](#)
- [Oracle Call Interface](#)
- [Custom Make Files](#)
- [Correcting Undefined Symbols](#)
- [Multi- threaded Applications](#)
- [XA Functionality](#)

## Overview of Oracle Precompilers

Oracle precompilers are application development tools used to combine SQL statements from an Oracle database with programs written in a high-level language. Oracle precompilers are compatible with ANSI SQL and are used to develop open, customized applications that run with Oracle8i or any other ANSI SQL database management system.

**See Also:** For general information on Oracle precompilers and interface features, see the *Oracle8i Server Application Developer's Guide*

## Precompiler Configuration Files

System configuration files for the Oracle Precompilers are located in the `$ORACLE_HOME/precomp/admin/` directory. [Table 4-1](#) lists the system configuration files.

**Table 4-1 System Configuration Files**

Product	Configuration File
Pro*C/C++ release 8.1.7	pcscfg.cfg
Pro*FORTRAN release 1.8.52	pccfor.cfg
Object Type Translator release 8.1.7	ottcfg.cfg

[Table 4-2](#) lists the location of precompiler README files. The README files describe changes since the last release.

**Table 4-2 Location of README Files for Oracle Products**

Product	README File
Pro*C/C++ release 8.1.7	<code>\$ORACLE_HOME/precomp/doc/proc2/readme.doc</code>
Pro*FORTRAN release 1.8.52	<code>\$ORACLE_HOME/precomp/doc/pro1x/readme.txt</code>

## Issues Common to All Precompilers

The following issues are common to all precompilers:

### Uppercase to Lowercase Conversion

In languages other than C, the compiler converts an uppercase function or subprogram name to lowercase. This can cause a “No such user exists” error message. If you receive this error message, verify that the function or subprogram name in your option file matches the case used in the `iapxtb` table.

### Vendor Debugger Programs

Precompilers and vendor-supplied debuggers can be incompatible. Oracle Corporation does not guarantee that a program run with a debugger will run the same way on an operating system.

### Value of `ireclen` and `oreclen`

The `ireclen` and `oreclen` parameters do not have maximum values.

## Static and Dynamic Linking

You can statically or dynamically link precompiler and OCI applications with Oracle libraries. With static linking, the libraries and objects of the whole application are linked together into a single executable program. As a result, application executables can become very large.

With dynamic linking, the executing code is partly stored in the executable program and partly stored in libraries that are linked dynamically by the application at run time. Libraries that are linked at run time are called dynamic or shared libraries.

Benefits of dynamic linking are:

- **Smaller disk requirements:** More than one application or invocation of the same application can use the same dynamic library.
- **Smaller main memory requirements:** The same dynamic library image is loaded into main memory only one time and can be shared by more than one application.

## Client Shared Library

The client shared library, `libclntsh.so`, is located in the `$ORACLE_HOME/lib` directory. If you use the Oracle provided `demo_product.mk` make file to link an application, the client shared library is used by default.

You may receive the following error message when starting an executable:

```
% sample1
ld.so.1: sample1: fatal: libclntsh.so.1.0: can't open file: errno=2
Killed
```

If you receive this error message, set the `LD_LIBRARY64_PATH` and `LD_LIBRARYN32_PATH` environment variables as follows:

```
$ LD_LIBRARY64_PATH=$ORACLE_HOME/lib: $ORACLE_HOME/javavm/admin
$ export LD_LIBRARY64_PATH
$ LD_LIBRARYN32_PATH=$ORACLE_HOME/lib32
$ export LD_LIBRARYN32_PATH
```

The client shared library is created automatically during installation. If you must re-create the client shared library:

1. Exit all client applications using the client shared library, including all Oracle client applications such as SQL\*Plus and Recovery Manager.
2. Log in as the `oracle` user and enter:

```
$ cd $ORACLE_HOME/rdbms/lib
$ make -f ins_rdbms.mk libclntsh.so
```

## Pro\*C/C++ Precompiler

Before you use Pro\*C/C++, verify that the MIPSpro C/C++ compiler release 7.2.1.3 or later is installed. This compiler is part of the MIPSpro compiler package supplied by SGI.

---

---

**Note:** To compile C source files containing multi-byte character references, you must use version 7.3 of the MIPSpro C compiler.

---

---

## Pro\*C/C++ Demonstration Programs

Demonstration programs are provided to show the varied functionality of the Pro\*C/C++ precompiler. There are three types of demonstration programs: C, C++, and Object programs. Object programs demonstrate the new Oracle8i Object features. All of the demonstration programs are located in the `$ORACLE_HOME/precomp/demo/proc` directory. The programs assume that the demonstration tables created by the `$ORACLE_HOME/sqlplus/demo/demobld.sql` script exist in the SCOTT schema with the password TIGER. By default, all programs are dynamically linked with the `$ORACLE_HOME/lib/libclntsh.so` client shared library.

**See Also:** For more information on using demonstration programs, see the *Pro\*C/C++ Precompiler Programmer's Guide*.

Use the `demo_proc.mk` make file, located in the `$ORACLE_HOME/precomp/demo/proc/` directory, to create the demonstration programs. For example, to precompile, compile, and link the `sample1` demonstration program, enter the following command:

```
$ make -f demo_proc.mk sample1
```

To create the C demonstration programs for Pro\*C/C++, enter:

```
$ make -f demo_proc.mk samples
```

To create the C++ demonstration programs for Pro\*C/C++, enter:

```
$ make -f demo_proc.mk cppsamples
```

To create the Object demonstration programs for Pro\*C/C++, enter:

```
$ make -f demo_proc.mk object_samples
```

Some demonstration programs require you to run a SQL script, located in the `$ORACLE_HOME/precomp/demo/sql` directory. If you do not run the script, a message displays requesting you to run it. To build a demonstration program and run the corresponding SQL script, include the `make` macro argument `RUNSQL=run` on the command line. For example, to create the `calldemo` demonstration program and run the required `$ORACLE_HOME/precomp/demo/sql/calldemo.sql` script, enter:

```
$ make -f demo_proc.mk calldemo RUNSQL=run
```

To create all Object demonstration programs and run all corresponding required SQL scripts, enter:

```
$ make -f demo_proc.mk object_samples RUNSQL=run
```

---

---

**Note:** The `ORA_CLIENT_LIB` environment variable is no longer used by the `demo_proc.mk` make file.

---

---

**See Also:** For information on using SQL\*Plus to build demonstration programs, see "[Using Demonstration Tables](#)" on page 3-3.

## Pro\*C/C++ User Programs

You can use the `$ORACLE_HOME/precomp/demo/proc/demo_proc.mk` make file to create programs. The syntax for linking a program with the `demo_proc.mk` make file is:

```
$ make -f demo_proc.mk target OBJS="objfile1 objfile2 ..." EXE=exename
```

In the preceding example:

- `target` is the make file target that you want to use (example, `build`)
- `objfilen` is the object file to link the program
- `EXE` is the executable program

For example, to create the program `myprog` from the Pro\*C/C++ source `myprog.pc`, enter one of the following commands, depending on the source and the type of executable you want to create.

For C source, dynamically linked with the client shared library, enter:

```
$ make -f demo_proc.mk build OBJS=myprog.o EXE=myprog
```

For C source, statically linked, enter:

```
$ make -f demo_proc.mk build_static OBJS=myprog.o EXE=myprog
```

For C++ source, dynamically linked with the client shared library, enter:

```
$ make -f demo_proc.mk cppbuild OBJS=myprog.o EXE=myprog
```

For C++ source, statically linked, enter:

```
$ make -f demo_proc.mk cppbuild_static OBJS=myprog.o EXE=myprog
```

## Pro\*FORTRAN Precompiler

Before using Pro\*FORTRAN, make sure that the MIPSpro FORTRAN compiler release 7.2.1.3 or later is installed. This compiler is part of the MIPSpro compiler package supplied by SGI.

## Pro\*FORTRAN Demonstration Programs

Demonstration programs are provided to show the varied functionality of the Pro\*FORTRAN precompiler. All of the demonstration programs are located in the `$ORACLE_HOME/precomp/demo/profor` directory and all of the programs assume that the demonstration tables created by the `$ORACLE_HOME/sqlplus/demo/demobld.sql` script exist in the SCOTT schema with the password TIGER. By default, all programs are dynamically linked with the client shared library `$ORACLE_HOME/lib/libclntsh.so`.

The makefile `demo_profor.mk` located in the `$ORACLE_HOME/precomp/demo/profor` directory should be used to create the demonstration programs.

For example, to precompile, compile, and link the `sample1` demonstration program, enter the command:

```
$ make -f demo_profor.mk sample1
```

Use the `$ORACLE_HOME/precomp/demo/profor/demo_profor.mk` makefile to create the demonstration programs. For example, to precompile, compile, and link the `sample1` demonstration program, enter one of the following commands:

```
$ make -f demo_profor.mk sample1
```

or

```
$ make -f demo_profor.mk build FORS=sample1.pfo EXE=sample1
```

To create all Pro\*FORTRAN demonstration programs, enter:

```
$ make -f demo_profor.mk samples
```

For some demonstration programs, you must run a SQL script, located in the `$ORACLE_HOME/precomp/demo/sql` directory. If you do not run the script, a message appears requesting you to run it. To build a demonstration program and run the corresponding SQL script, include the make macro argument `RUNSQL=run` on the command line. For example, to create the `sample11` demonstration program and run the required `$ORACLE_HOME/precomp/demo/sql/sample11.sql` script, enter:

```
$ make -f demo_profor.mk sample11 RUNSQL=run
```

**See Also:** For information on using SQL\*Plus to build demonstration programs, see ["Using Demonstration Tables"](#) on page 3-3. For more information on Pro\*FORTRAN precompilers, see *Pro\*FORTRAN Supplement to Oracle Precompilers*.

## Pro\*FORTRAN User Programs

You can use the `$ORACLE_HOME/precomp/demo/profor/demo_profor.mk` makefile to create user programs. The syntax for linking a user program with `demo_profor.mk` is:

```
$ make -f demo_profor.mk target FORS="forfile1 forfile2 ..." EXE=exename
```

In the preceding example:

- `target` is the makefile target that you want to use, for example `build`
- `objfilen` is the object file to link the program
- `exename` is the executable program

For example, to create the program, `myprog`, from the Pro\*FORTRAN source `myprog.pfo`, enter one of the following commands, depending on the type of executable that you want:

For a dynamically linked executable with the client shared library, enter:

```
$ make -f demo_profor.mk build FORS=myprog.f EXE=myprog
```

For a statically linked executable, enter:

```
$ make -f demo_profor.mk build_static FORS=myprog.f EXE=myprog
```

## Oracle Call Interface

Before using the Oracle Call Interface (OCI), make sure that the MIPSpro C/C++ compiler release 7.2.1.3 or later is installed. This compiler is part of the MIPSpro compiler package supplied by SGI.

**See Also:** The required version for your operating system is specified in the *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX*.

## OCI Demonstration Programs

Demonstration programs are provided to show the varied functionality of the OCI. There are two types of demonstration programs: C and C++. All of the demonstration programs are located in the `$ORACLE_HOME/rdbms/demo` directory. By default, all programs are dynamically linked with the client shared library `$ORACLE_HOME/lib/libclntsh.so`.

Many of the demonstration programs assume that the demonstration tables created by the `$ORACLE_HOME/sqlplus/demo/demobld.sql` script exist in the SCOTT schema with the password TIGER.

Use the `$ORACLE_HOME/rdbms/demo/demo_rdbms.mk` make file to create the demonstration programs. For example, to compile and link the `cdemo1` demonstration program, enter:

```
$ make -f demo_rdbms.mk cdemo1
```

To create the OCI C demonstration programs, enter:

```
$ make -f demo_rdbms.mk demos
```

To create all of the OCI C++ demonstration programs, enter:

```
$ make -f demo_rdbms.mk cppdemos
```

Some demonstration programs require you to run a SQL script found in the `$ORACLE_HOME/rdbms/demo` directory. If you do not run the script, a message appears requesting you to run the script. In most cases, the SQL script name is the same as the program name with a `.sql` extension.

**See Also:** For information on using SQL\*Plus to build demonstration programs, see "[Using Demonstration Tables](#)" on page 3-3.

For further information on the demonstration programs see the *Programmer's Guide to the Oracle Call Interface* and the program source for details of each program.

## OCI User Programs

You can use the `demo_rdbms.mk` make file, located in the `$ORACLE_HOME/rdbms/demo` directory, to create programs. The syntax for the linking a user program with `demo_rdbms.mk` is:

```
$ make -f demo_rdbms.mk target OBJS="objfile1 objfile2 ..." EXE=exename
```

In the preceding example:

- `target` is the make file target that is used (example, build)
- `objfilen` is the object file to link the program
- `EXE` is the executable program

For example, to create the `myprog` program from the C/C++ source `myprog.c`, enter one of the following, depending on the type of executable you want:

To create the `myprog` program from the C source, dynamically linked with a the client shared library, enter:

```
$ make -f demo_rdbms.mk build OBJS=myprog.o EXE=myprog
```

To create the `myprog` program from the C source, statically linked, enter:

```
$ make -f demo_rdbms.mk build_static OBJS=myprog.o EXE=myprog
```

To create the `myprog` program from the C++ source, dynamically linked with client shared library, enter:

```
$ make -f demo_rdbms.mk buildc++ OBJS=myprog.o EXE=myprog
```

To create the `myprog` program from the C++ source, statically linked, enter:

```
$ make -f demo_rdbms.mk buildc++_static OBJS=myprog.o EXE=myprog
```

---

---

**Note:** In the preceding examples, the file `myprog.o` is the object file generated by the compiler.

---

---

## Custom Make Files

Oracle Corporation recommends that you use the provided `demo_product.mk` make files to link user programs as described in the specific product sections of this chapter. If you modify the provided make file, or if you choose to use a custom-written make file, the following restrictions apply:

- Do not modify the order of the Oracle libraries. Oracle libraries are included on the link line more than once so that all the symbols are resolved during linking. The order of the Oracle libraries is essential for the following reasons:
  - Oracle libraries are mutually referential. Functions in library A call functions in library B, and functions in library B call functions in library A.
  - The SGI IRIX linker is a one-pass linker. The linker searches a library exactly once at the point it is encountered in the link line.
- If you add your own library to the link line, add it to the beginning or to the end of the link line. Do not place user libraries between the Oracle libraries.
- If you choose to use a make utility such as `nmake` or GNU `make`, you should be aware of how macro and suffix processing differs from the `make` utility provided with SGI IRIX. Oracle make files are tested and are supported with the SGI IRIX `make` utility.
- Oracle library names and the contents of Oracle libraries are subject to change between releases. Always use the `demo_product.mk` make file that ships with the current release as a guide to determine the required libraries.

## Correcting Undefined Symbols

Oracle provides the `symfind` utility to assist in locating a library or object file where a symbol is defined. A common error when linking a program is undefined symbols, which produces an error message similar to the following:

```
$ make -f demo_proc.mk sample1
Undefined                          first referenced
  symbol                            in file
sqlcex                             sample1.o
sqlglm                              sample1.o
ld: fatal: Symbol referencing errors. No output written to sample1
```



Table 4–3 describes the signals that Oracle8i uses for two-task communications.

**Table 4–3 Signals for Two-Task Communications**

Signal	Description
SIGCLD	The pipe driver uses SIGCLD, also referred to as SIGCHLD, when an Oracle process dies. The UNIX kernel sends a SIGCLD signal to the user process. The signal handler uses the wait() routine to see if a server process died. The Oracle process does not catch SIGCLD; the user process catches it.
SIGCONT	The pipe two-task driver uses SIGCONT to send out-of-band breaks from the user process to the Oracle process.
SIGINT	Two-task drivers use SIGINT to detect user interrupt requests. The Oracle process does not catch SIGINT; the user process catches it.
SIGIO	Net8 protocols use SIGIO to indicate incoming networking events.
SIGPIPE	The pipe driver uses SIGPIPE to detect end-of-file on the communications channel. When writing to the pipe, if no reading process exists, a SIGPIPE signal is sent to the writing process. Both the Oracle process and the user process catch SIGPIPE. SIGCLD is similar to SIGPIPE, but only applies to user processes, not to Oracle processes.
SIGTERM	The pipe driver uses SIGTERM to signal interrupts from the user to the Oracle process. This occurs when the user presses the interrupt key [Ctrl]+[c]. The user process does not catch SIGTERM; the Oracle process catches it.
SIGURG	Net8 TCP/IP drivers use SIGURG to send out-of-band breaks from the user process to the Oracle process.

The listed signals affect Pro\*C or other precompiler applications. You can install one signal handler for SIGCLD (or SIGCHLD) and SIGPIPE when connected to the Oracle process. If you call the `osnsui()` routine to set it up, you can have more than one signal handle for SIGINT. For SIGINT, use `osnsui()` and `osncui()` to register and delete signal-catching routines. You can also install as many signal handlers as you want for other signals. If you are not connected to the Oracle process, you can have multiple signal handlers.

The following example shows how you can set up a signal routine and a catching routine:

```
/* user side interrupt set */
word osnsui( /*_ word *handlp, void (*astp), char * ctx, _*/)
/*
** osnsui: Operating System dependent Network Set User-side Interrupt. Add an
** interrupt handling procedure astp. Whenever a user interrupt(such as a ^C)
** occurs, call astp with argument ctx. Put in *handlp handle for this
** handler so that it may be cleared with osncui. Note that there may be many
** handlers; each should be cleared using osncui. An error code is returned if
** an error occurs.
*/

/* user side interrupt clear */
word osncui( /*_ word handle _*/ );
/*
** osncui: Operating System dependent Clear User-side Interrupt. Clear the
** specified handler. The argument is the handle obtained from osnsui. An error
** code is returned if an error occurs.
*/
```

The following example is a template for using the `osnsui()` and the `osncui()` routines in an application program:

```
/*
** User interrupt handler template.
*/
void sig_handler()
{
...
}

main(argc, argv)
int arc;
char **argv;
{

    int handle, err;
    ...

    /* set up my user interrupt handler */
    if (err = osnsui(&handle, sig_handler, (char *) 0))
    {
        /* if the return value is non-zero, an error has occurred
        Take appropriate action for the error. */
    }
}
```

```
...
}
...
/* clear my interrupt handler */
if (err = osncui(handle))
{
/* if the return value is non-zero, an error has occurred
Take appropriate action for the error. */
...
}
...
}
```

## XA Functionality

Oracle XA is the Oracle implementation of the X/Open Distributed Transaction Processing (DTP) XA interface. The XA standard specifies a bi-directional interface between resource managers (for example, Oracle) that provide access to shared resources within transactions, and between a transaction service that monitors and resolves transactions.

Oracle Call Interface has XA functionality. When building a TP-monitor XA application, ensure that the TP-monitor libraries (that define the symbols `ax_reg` and `ax_unreg`) are placed in the link line before the Oracle client shared library. This link restriction is required only when using the XA dynamic registration (Oracle XA switch `xaoswd`).

Oracle8i does not support Oracle7 release 7.1.6 XA calls (although it does support release Oracle7 release 7.3 XA calls). Therefore, TP-monitor XA applications using Oracle7 release 7.1.6 XA calls must be relinked with the Oracle8i XA library. The Oracle8i XA calls are defined in both the `$ORACLE_HOME/lib/libclntsh.so` shared library and the `$ORACLE_HOME/lib/libclient8.a` static library.



---

---

# Configuring Net8

This chapter describes Net8 features on Oracle8i release 3 (8.1.7) 64 Bit for SGI IRIX. It includes the following sections:

- [Core Net8 Products and Features](#)
- [Net8 Protocol Support](#)
- [BEQ Protocol](#)
- [IPC Protocol](#)
- [RAW Protocol](#)
- [TCP/IP Protocol Adapter](#)
- [Configuring Oracle Intelligent Agent for Oracle SNMP](#)
- [Oracle Advanced Security](#)

## Core Net8 Products and Features

This section describes core Net8 products and features.

### Net8 Files and Utilities

This section describes the files and utilities that you can use to configure Oracle Net8 products.

#### Location of Net8 Configuration Files

The default directory for Net8 configuration files on SGI IRIX systems is `/etc`.

Oracle Net8 searches the following locations for global files in the following order:

1. The directory specified by the `TNS_ADMIN` environment variable, if set
2. The `/etc` directory
3. The `$ORACLE_HOME/network/admin` directory

To set the `TNS_ADMIN` environment variable in the startup files of all network users, enter the following for the Bourne or Korn shell:

```
$ TNS_ADMIN=directory_path
$ export TNS_ADMIN
```

For each system level configuration file, users may have a corresponding local private configuration file (stored in the user's home directory). The settings in the local file override the settings in the system level file. The local configuration file for `sqlnet.ora` is `$HOME/.sqlnet.ora`. The local configuration file for `tnsnames.ora` is `$HOME/.tnsnames.ora`. Syntax for these files is identical to that of the corresponding system files.

#### Sample Configuration Files

Examples of the `cman.ora`, `listner.ora`, `names.ora`, `sqlnet.ora`, and `tnsnames.ora` configuration files are located in the `$ORACLE_HOME/network/admin/samples` directory.

## The Adapters Utility

Net8 provides support for various network protocol adapters and naming methods. These adapters and naming methods are linked to particular executables and provide the interface between network protocols and Net8.

To display installed Net8 protocol adapters, enter:

```
$ adapters
```

To display adapters linked with a specific executable, enter:

```
$ adapters executable
```

For example, to display the Net8 protocol adapters linked with the Oracle executable, enter:

```
$ adapters oracle
```

The `adapters` executable displays output similar to the following:

```
Net8 Protocol Adapters linked with oracle are:
```

```
BEQ Protocol Adapter
```

```
IPC Protocol Adapter
```

```
TCP/IP Protocol Adapter
```

```
RAW Protocol Adapter
```

```
Net8 Naming Adapters linked with oracle are:
```

```
Oracle TNS Naming Adapter
```

```
Oracle Naming Adapter
```

```
Oracle Advanced Security/Networking Security products linked with oracle are:
```

## Oracle Connection Manager

Oracle Connection Manager (OCM) is a router through which a client connection request may be sent either to its next hop or directly to the database server. Clients who route their connection requests through the OCM can then take advantage of the connection concentration, Net8 access control, or multi-protocol support features configured on that OCM.

OCM listens for incoming requests from clients and initiates connection requests to destination services. OCM performs these tasks with the help of Oracle Connection Gateway Process and Oracle Connection Manager Administrative Process.

OCM also includes a feature which you can use to control client access to designated servers in a TCP/IP environment. By specifying certain filtering rules you may allow or restrict access from specific clients to a server based on the following criteria:

- Source host names or IP addresses for clients
- Destination host names or IP addresses for servers
- Destination database service name

### **Net8 Firewall Proxy**

Some firewall vendors also offer Net8 Firewall Proxy, which is installed on firewalls requiring an application proxy. Net8 Firewall Proxy has the same access controls as Oracle Connection Manager.

**See Also:** For more information on the Oracle Connection Manager, see the *Net8 Administrator's Guide*.

## **Server Models**

Net8 connects clients and servers using two different models:

- **Multi-threaded Server Model**—The listener passes the connection to the dispatcher or redirects the clients to connect through a dispatcher.

In a Multi-threaded Server (MTS) configuration, client user processes connect to a dispatcher. A dispatcher can support multiple client connections concurrently. Each client connection is bound to a virtual circuit. A virtual circuit is a piece of shared memory used by the dispatcher for client database connection requests and replies. This approach enables a small pool of server processes to serve a large number of clients. A significant advantage of the MTS model over the dedicated server model is the reduction in the use of system resources, enabling the support of an increased number of users.

- **Dedicated Server Model**—The listener starts a dedicated server and passes the connection to the dedicated server or redirects the clients to the dedicated server.

In the dedicated server model, there is one server process for each client. For clients to connect to dedicated servers, the listener and the instance must be running on the same system. Dedicated server processes require more memory than MTS.

**See Also:** For more information on the Multi-threaded Server, see the *Net8 Administrator's Guide*.

## Oracle Names

Oracle Names uses Oracle proprietary software to store the names and addresses of all database services on a network. Clients that want to connect to a server direct their connection requests to an Oracle Names server. Oracle Names servers resolve the name to a network address and return that information to the client.

**See Also:** For information on Oracle Names, see the *Net8 Administrator's Guide*.

## Net8 Protocol Support

Net8 release 8.1.7 on SGI IRIX supports the following protocols:

- BEQ
- IPC
- RAW
- TCP/IP

Before installing the TCP/IP protocol, you must install the appropriate operating system software. The BEQ, IPC, and RAW Net8 protocols do not have any specific operating system requirements.

**See Also:** Refer to the *Oracle8i Installation Guide Release 3 (8.1.7) 64 Bit for SGI IRIX* for more information on Net8.

## ADDRESS Specification

The IPC and TCP/IP Net8 protocols each have a protocol-specific ADDRESS specification that is used for Net8 configuration files and for the MTS\_LISTENER\_ADDRESS database initialization parameter in the `initsid.ora` file. See the ADDRESS specification heading under each protocol section in this chapter for details.

Table 5–1 shows a summary of ADDRESS specifications for each protocol.

**Table 5–1 ADDRESS Specification Summary**

Supported Protocol	ADDRESS Specification
BEQ	(ADDRESS = (PROTOCOL = BEQ) (PROGRAM = ORACLE_HOME/bin/oracle) (ARGV0 = oracleORACLE_SID) (ARGS = '(DESCRIPTION=(LOCAL=YES)(ADDRESS=(PROTOCOL=BEQ)))') (ENVS = 'ORACLE_HOME=ORACLE_HOME,ORACLE_SID=ORACLE_SID') )
IPC	(ADDRESS = (PROTOCOL=IPC) (KEY=key) )
RAW	N/A
TCP/IP	(ADDRESS = (PROTOCOL=TCP) (HOST=hostname) (PORT=port_id) )

## BEQ Protocol

The BEQ protocol adapter is both a communications mechanism and a process-spawning mechanism. To use the BEQ protocol adapter, the client and server must be on the same system. If a network service name is not specified, either directly by the user on the command line or Login screen or indirectly through an environment variable such as TWO\_TASK, then the BEQ protocol is used. In this case, the BEQ protocol always uses a dedicated server and the multi-threaded server is never used. This dedicated server is started automatically by the BEQ protocol, which waits for the server process to start and attach to an existing System Global Area (SGA). If the startup of the server process is successful, the BEQ protocol then provides inter-process communication through UNIX pipes.

An important feature of the BEQ protocol is that no listener is required for its operation. The protocol is linked into the client tools and directly starts its own server process without outside interaction. However, the BEQ protocol can only be used when the client program and Oracle8i are installed on the same system. The BEQ protocol is always installed and always linked to all client tools and to the Oracle8i server.

The BEQ protocol connection parameters are part of the ADDRESS keyword-value pair. The ADDRESS is commonly part of a larger construct such as a connect descriptor or configuration file. You can enter the following parameters in any order:

```
(ADDRESS =
  (PROTOCOL = BEQ)
  (PROGRAM = ORACLE_HOME/bin/oracle)
  (ARGV0 = oracleORACLE_SID)
  (ARGS = '(DESCRIPTION=(LOCAL=YES)(ADDRESS=(PROTOCOL=BEQ)))')
  (ENVS = 'ORACLE_HOME=ORACLE_HOME,ORACLE_SID=ORACLE_SID')
)
```

Table 5–2 describes the syntax for the BEQ protocol connection parameters.

**Table 5–2 Syntax for BEQ Protocol Connection Parameters**

Parameter	Description
PROTOCOL	The protocol adapter to be used. The value is BEQ. It is not case sensitive.
PROGRAM	The full path to the <code>oracle</code> executable.
ARGV0	The name of the process as it appears in a <code>ps</code> listing. The recommended value is <code>oracleORACLE_SID</code> .
ARGS	<code>'(DESCRIPTION=(LOCAL=YES)(ADDRESS=(PROTOCOL=BEQ)))'</code>
ENVS	Environment specification that includes the <code>ORACLE_HOME</code> and <code>ORACLE_SID</code> variables, where <code>ORACLE_HOME</code> represents the full path to the Oracle home directory of the database, and <code>ORACLE_SID</code> represents the system identifier of the database.

**Example 5–1** shows a sample BEQ ADDRESS.:

**Example 5–1 BEQ ADDRESS Specifying a Client**

```
(ADDRESS =  
  (PROTOCOL = BEQ)  
  (PROGRAM = /u01/app/oracle/product/8.1.7/bin/oracle)  
  (ARGV0 = oracleV817)  
  (ARGS = '(DESCRIPTION=(LOCAL=YES)(ADDRESS=(PROTOCOL=BEQ)))')  
  (ENVS = 'ORACLE_HOME=/u01/app/oracle/product/8.1.7,ORACLE_SID=V817')
```

## IPC Protocol

The IPC protocol is similar to the BEQ protocol in that it can only be used when the client program and the Oracle8i server are installed on the same system. The IPC protocol differs from the BEQ protocol in that it can be used both with a dedicated server and Multi-threaded Server configurations. The IPC protocol requires a listener for its operation. The IPC protocol is always installed and always linked to all client tools and to Oracle8i.

For the IPC protocol, the location of the UNIX Domain Socket (IPC) file on UNIX systems changed after Oracle7 release 7.1. If you have Oracle7 release 7.1 installed on the same computer as Oracle8i and you attempt to make an IPC connection between the two instances, the connection might fail. The solution to this problem is to make a symbolic link between the previous location of the IPC file directory (`/var/tmp/o`) to its current directory location (`/var/tmp/.oracle`).

The IPC protocol connection parameters are part of the ADDRESS keyword-value pair. The ADDRESS is commonly part of a larger construct such as a connect descriptor or configuration file. You can enter the following parameters in any order:

```
(ADDRESS=  
  (PROTOCOL=IPC)  
  (KEY=key)  
)
```

[Table 5–3](#) describes the syntax for IPC protocol connections parameters.

**Table 5–3 Syntax for IPC Protocol Connection Parameters**

Parameter	Description
PROTOCOL	IPC protocol adapter to be used. The value is IPC. It is not case sensitive.
KEY	Service name of database or database identifier (SID).

[Example 5–2](#) shows a sample IPC ADDRESS.

**Example 5–2 IPC Address Specifying a Client**

```
(ADDRESS=
  (PROTOCOL=IPC)
  (KEY=PROD)
```

The ADDRESS is commonly part of a larger construct such as a connect descriptor or configuration file.

## RAW Protocol

When data is transferred between a client and a server, Net8 adds its own header information to every network packet. Through the Raw Transport feature, Net8 can now minimize header information on each packet going over the network.

After a connection is established, two types of information flow over the network: data and break handling. The connection packets require the Net8 header information to establish the connection correctly. However, after the connection is established, all data packets are stripped of their Net8 header information and passed directly to the operating system, bypassing the Net8 network and protocol layers. The performance of the connection is increased because of fewer protocol stack layers for the data to flow through and fewer bytes that are transmitted over the network.

This feature is transparently enabled when it is required. If no existing features require that header information be transmitted, the headers are stripped off. For example, Raw Transport would not be enabled when you use authentication, which requires certain information to be sent with each packet of information.

This feature requires no configuration. Net8 determines whether the conditions are met and then transparently switches to Raw Transport mode.

## TCP/IP Protocol Adapter

Oracle Corporation recommends that you reserve a port for your listener in the `/etc/services` file of each Net8 node on the network. The default port number is 1521. The entry lists the listener name and the port number, for example:

```
listener      1521/tcp
```

In this example *listener* is the name of the listener, as defined in the `listener.ora` file. You should reserve more than one port to start more than one listener.

## Specifying a TCP/IP ADDRESS

The TCP/IP protocol connection parameters are part of the ADDRESS keyword-value pair. The ADDRESS is commonly part of a larger construct such as a connect descriptor or configuration file. You can enter the parameters in any order:

```
(ADDRESS=  
  (PROTOCOL=TCP)  
  (HOST=hostname)  
  (PORT=port_id)  
)
```

[Table 5–4](#) describes the syntax for the TCP/IP protocol connection parameters.

**Table 5–4 Syntax for TCP/IP Protocol Connection Parameters**

Parameter	Description
PROTOCOL	The protocol adapter to be used. The value is TCP. It is not case sensitive.
HOST	The host name or the host IP address.
PORT	The TCP/IP port. Either a number or the name specified in the <code>/etc/services</code> file. Oracle Corporation recommends a value of 1521.

[Example 5–3](#) shows a sample TCP/IP ADDRESS.

**Example 5-3 TCP/IP ADDRESS Specifying a Client**

```
(ADDRESS=
  (PROTOCOL=TCP)
  (HOST=MADRID)
  (PORT=1521)
)
```

You can specify the last field by name, for example, (PORT=listener).

## Configuring Oracle Intelligent Agent for Oracle SNMP

Although Oracle Intelligent Agent does not require Simple Network Management Protocol (SNMP), you can configure Oracle SNMP support before starting the Intelligent Agent. Note that all of the configuration files for the following steps are located in the `$ORACLE_HOME/network/snmp/peer` directory.

### Configure the Encapsulator

To configure the encapsulator, perform the following steps:

1. Add the following line to the `snmpd.conf` file, where *hostname\_or\_IP\_address* represents the local system IP address:

```
trap hostname_or_IP_address
```

2. In the `CONFIG.encap` file, you can modify the port number which is set to 1161 in the default file. If you modify the port number, you must also modify the port number for `NEW_SNMPD_PORT` in the `start_peer` script.

`NEW_SNMPD_PORT` is the port on which the `snmpd` agent (the native SGI IRIX SNMP agent) listens. Make sure that this is the same port as specified in the `CONFIG.encap` file. `NEW_TRAPD_PORT` is the PEER encapsulator port to which the `snmpd` agent sends traps.

`NEW_SNMPD_PORT` and `NEW_TRAPD_PORT` in the `start_peer` script must have different port numbers. You can also modify the `NEW_TRAPD_PORT` port number.

### Verify start\_peer Script

The `start_peer` script contains a line similar to the following:

```
SNMPD = snmpd_executable_path
```

If the `snmpd` executable on your system is not in the location indicated by this line, specify the correct location of the `snmpd` executable.

### Start the SNMP Components

To start the SNMP components, perform the following steps:

1. Verify that the SNMP components (`master_peer`, `encap_peer`, and `snmpd`), are *not* running:

```
$ ps -aef | grep peer
$ ps -aef | grep snmp
```

If any of the components are running, log in as the `root` user and use the `kill` command to terminate the processes before proceeding.

2. As the `root` user, run the `start_peer` script to start the PEER master agent, PEER encapsulator, and native SGI IRIX SNMP agent:

```
# cd $ORACLE_HOME/network/snmp/peer
# ./start_peer -a
```

---

---

**Note:** If you do not have the native SGI IRIX SNMP agent in your system, you must *not* use the PEER encapsulator. To start the master agent only, run `start_peer -m`.

---

---

3. Verify that the SNMP components are running:

```
# ps -aef | grep peer
# ps -aef | grep snmp
```

### Configure and Start the Database Subagent

The configuration and startup of the database subagent (the Oracle Intelligent Agent) is described in the *Oracle Enterprise Manager Configuration Guide*.

## Oracle Advanced Security

The following section describes the Oracle Advanced Security Option.

### **.bak Files**

When you install Oracle Advanced Security, three .bak files are created: naeet.o.bak, naect.o.bak, and naedhs.o.bak. They are located in the \$ORACLE\_HOME/lib directory. These files are required for relinking during Oracle Advanced Security de-installation and should not be deleted.

### **Security and Single Sign-On**

For more information on configuring Security and Single Sign-On, see the *Oracle Advanced Security Administrator's Guide*.



---

---

# Running Oracle Data Option Demonstrations

This chapter describes Oracle Data Option demonstrations. It contains the following sections:

- [Oracle8i interMedia](#)
- [Oracle8i Time Series Demonstrations](#)
- [Oracle8i Visual Information Retrieval](#)
- [Oracle8i Spatial](#)

## Oracle8i *interMedia*

Oracle8i *interMedia* includes the following components:

- Text
- Audio, Video, and Image
- Locator
- Web Agent and Clipboard

### Text

There are no demonstrations for *interMedia* Text in Oracle8i. However, *interMedia* Text now includes code samples. Use a web browser to open the following file:

```
$ORACLE_HOME/ctx/sample/api/index.html
```

**See Also:** For more information on *intermedia* Text, see the *Oracle8i interMedia Text Reference*, and the *Oracle8i ConText interMedia Text Migration* guide.

### Audio, Video, and Image

Oracle8i *interMedia* includes audio, video, and image scripts and sample programs in the following directories:

```
$ORACLE_HOME/ord/aud/demo/  
$ORACLE_HOME/ord/img/demo/  
$ORACLE_HOME/ord/vid/demo/
```

#### Sample Audio Scripts

The *interMedia* audio scripts consist of the following files:

- `auddemo.sql` - demonstration that shows features of the audio object including:
  - Checking *interMedia* objects
  - Creating a sample table containing audio
  - Inserting NULL rows into the audio table
  - Checking the rows out
  - Checking all the audio attributes directly

- Checking all the audio attributes by calling methods
- Installing your own format plug-in using the `fplugins.sql` and `fpluginb.sql` files
- `fplugins.sql` - demonstration format plug-in specification that you can use as a guideline to write any format plug-in you want to support.
- `fpluginb.sql` - demonstration format plug-in body that you can use as a guideline to write any format plug-in you want to support.

See the `README.txt` file in the `$ORACLE_HOME/ord/aud/demo` directory for requirements and instructions on running this SQL demonstration.

### Sample Program for Modifying Images or Testing the Image Installation

After you have installed Oracle8i *interMedia* Image, you can run the Oracle8i *interMedia* Image demonstration program. You can also use this program to confirm successful installation.

This section lists the steps required to build and run the *interMedia* Image demonstration.

The *interMedia* Image demonstration files are located in the `$ORACLE_HOME/ord/img/demo` directory, where `$ORACLE_HOME` is the Oracle home directory.

### Demonstration Installation Steps

To install the Oracle8i *interMedia* Image demonstration:

1. Create the SCOTT/TIGER database user, if this user does not exist:

```
$ svrmgrl
SVRMGR> CONNECT INTERNAL;
SVRMGR> CREATE USER SCOTT IDENTIFIED BY TIGER;
SVRMGR> GRANT CONNECT,RESOURCE TO SCOTT;
```

2. Create the image demonstration directory in the Oracle home directory:

```
$ svrmgrl
SVRMGR> CONNECT INTERNAL;
SVRMGR> CREATE OR REPLACE DIRECTORY IMGDEMODIR AS '$ORACLE_
HOME/ORD/IMG/DEMO';
```

3. Grant privileges on the directory to PUBLIC:

```
SVRMGR> grant read on directory imgdemodir to public with grant option;
```

4. If required, make the `imgdemo` program:

```
$ cd $ORACLE_HOME/ord/img/demo
$ make -f demo_ordimg.mk imgdemo
```

### Running the Demonstration

The `imgdemo` file is a sample program that shows how Oracle8i *interMedia* Image can be used from within a program. The demonstration is written in C and uses the Oracle Call Interface (OCI) to access the database and exercise Oracle8i *interMedia* Image.

The program operates on the `imgdemo.dat` file, which is a bitmap (BMP) image in the demonstration directory. Optionally, you can supply the name of an image file on the command line, provided that the file is located in the same directory as the demonstration. In either way, after the image has been manipulated by Oracle8i *interMedia* Image, the resulting image is written to the `imgdemo.out` file. You can then view the image with an image viewer such as `xv`.

When you run the demonstration, it deletes and re-creates a table named `IMGDEMOTAB` in the `SCOTT/TIGER` schema of the default database. This table is used to hold the demonstration data. After the table is created, a reference to the image file is inserted into the table. The data is then loaded into the table and converted to the JFIF format using the `processCopy()` method of `ORDImage`.

The demonstration extracts the image properties from the database using the `setProperties()` method. After the demonstration runs `setProperties()` it issues an `UPDATE` command. This is required because `setProperties()` only updates a local copy of the type attributes.

Next, the demonstration uses the Oracle8i *interMedia* Image `process()` method to cut and scale the image within the database. The demonstration then runs an update that commits the change. The demonstration cuts a portion of the image 100 pixels wide by 100 pixels high starting from pixel location (100,100). This sub-image is scaled to twice its original size and the resulting image is written to a file named `imgdemo.out` in the current directory.

## Executing the Demonstration from the Command Line

To execute the demonstration, enter:

```
$ imgdemo [image-file]
```

The optional image file must be located in the same directory as the demonstration program. The demonstration displays a number of messages describing its progress, along with any errors it encounters. Expect to see the following messages:

```
Dropping table IMGDEMOTAB...
Creating and populating table IMGDEMOTAB...
Loading data into cartridge...
Modifying image characteristics...
Writing image to file imgdemo.out...
Disconnecting from database...
Logged off and detached from server.
Demo completed successfully.
```

If the program encounters any errors, it is likely that either Oracle8i *interMedia* software has not been installed correctly or the database has not been started. If the program completes successfully, you can view the original image and the resultant image which has undergone the cutting and scaling described previously, with an image viewer such as `xv`.

## Sample Video Scripts

The *interMedia* Video scripts consist of the following files:

- `viddemo.sql` - demonstration that shows features of the video object including:
  - Checking *interMedia* objects
  - Creating a sample table with video in it
  - Inserting NULL rows into the video table
  - Checking the rows out
  - Checking all the video attributes directly
  - Checking all the video attributes by calling methods
  - Installing your own format plug-in using the two files, `fplugins.sql` and `fpluginb.sql`

- `fplugins.sql` - demonstration format plug-in specification that you can use as a guideline to write any format plug-in you want to support.
- `fpluginb.sql` - demonstration format plug-in body that you can use as a guideline to write any format plug-in you want to support

**See Also:** See the `README.txt` file in the `$ORACLE_HOME/ord/vid/demo` directory for requirements and instructions on running this SQL demonstration.

### Java Demonstration

Oracle8i *interMedia* Locator includes a Java demonstration to help you learn to use both the audio and video client-side Java classes for building your own applications. These two demonstrations instantiate the audio and video object at the client side and start a number of accessor methods. The audio Java demonstration files are located in the `$ORACLE_HOME/ord/aud/demo` directory and the video Java demonstration files are located in the `$ORACLE_HOME/ord/vid/demo` directory.

**See Also:** See the `README.txt` file in each directory for requirements and instructions on running each Java demonstration.

### MediaAnnotator

The MediaAnnotator program is not contained on the Oracle8i *interMedia* CD. It (along with other free Oracle software) is available at the following URL:

[http://technet.oracle.com/software/products/intermedia/software\\_index.htm](http://technet.oracle.com/software/products/intermedia/software_index.htm)

**See Also:** See the *Oracle8i interMedia Audio, Image, and Video User's Guide and Reference* and *Oracle8i interMedia Audio, Image, and Video Java Client User's Guide and Reference* for more information on MediaAnnotator.

## Locator

Oracle8i *interMedia* Locator includes a number of scripts that you can modify and run.

**See Also:** See the *Oracle8i interMedia Locator User's Guide and Reference* for more information on *interMedia* Locator.

## Sample Scripts

Sample Oracle8i *interMedia* Locator scripts are available in the following directory after you install *interMedia* Locator:

```
$ORACLE_HOME/md/demo/geocoder
```

These scripts consist of the following files:

- `geohttp.sql`

This file contains two parts that use different modes to run the geocode function. These modes are:

- Interactive mode

See the *Oracle8i interMedia Locator User's Guide and Reference* for a listing of this part of the file.

- Batch mode

You must update the setup tables in the `nh_cs.sql` file before you run the `geohttp.sql` in batch mode. See the *Oracle8i interMedia Locator User's Guide and Reference* for a listing of this part of the file.

- `geoindex.sql`

This file contains:

- A function named `ESTIMATE_LEVEL` to better estimate the index level for use with the spatial locator index for within-distance queries that use a radius distance greater than 100 miles. For a listing of this file, see the *Oracle8i interMedia Locator User's Guide and Reference*.

- A procedure statement named `SETUP_LOCATOR_INDEX` that builds a setup spatial locator index on the location column that contains the spatial information within the `CUST_TABLE` table where the spatial information is stored. For a listing of this file, see the *Oracle8i interMedia Locator User's Guide and Reference*.

- `geolocate.sql`

This file contains a routine that dynamically creates a geometry of interest and then queries against the `NH_COMPUTER_STORES` table to find out how many stores are within a 10-mile radius of the office. For a listing of this file, see the *Oracle8i interMedia Locator User's Guide and Reference*.

## Web Agent and Clipboard

The Web Agent and Clipboard are not contained on the Oracle8i *intermedia* CD. They (along with other free Oracle software) are available at the following URL:

[http://technet.oracle.com/software/products/intermedia/software\\_index.htm](http://technet.oracle.com/software/products/intermedia/software_index.htm)

**See Also:** See *Using Oracle8i interMedia with the Web* for more information on Web Agent.

## Oracle8i Time Series Demonstrations

Table 6–1 shows the demonstrations included with Oracle8i Time Series. This table includes a description of each demonstration and the default directory in which its files are installed.

In the following table, all directories listed are sub directories of the \$ORACLE\_HOME/ord/ts/demo directory:

**Table 6–1 Oracle8i Time Series Demonstrations**

Description	Sub Directory
Quick-start demonstration: quick start using Oracle8i Time Series. (See the <i>Oracle8i Time Series User's Guide</i> for more information).	tsquick
Usage demonstration for end users and product developers who want to use existing Oracle8i Time Series features. (See the <i>Oracle8i Time Series User's Guide</i> ).	usage
Electric utility application demonstrating how to calculate peak and off-peak summaries of 15-minute data.	usageutl
Java-based retrieval of time series data, using the prototype Oracle8i Time Series Java API and designed to run in a Web browser. (See the <i>Oracle8i Time Series User's Guide</i> for more information).	applet
Simple Java code segments that perform time series operations and print the results. (See the <i>Oracle8i Time Series User's Guide</i> for more information).	java
Demonstration showing the use of administrative tools procedures to "retrofit" existing time series detail tables. Demonstration showing how to support time series queries for multiple qualifier columns in the time series detail table.	retrofit
Advanced-developer demonstration for those who want to extend Oracle8i Time Series features.	extend

**Table 6–1 Oracle8i Time Series Demonstrations (Cont.)**

Description	Sub Directory
OCI demonstration showing how to call Oracle8i Time Series functions using the Oracle Call Interface.	oci
PRO*C/C++ demonstration showing how to call Oracle8i Time Series functions in applications created using the Oracle Pro*C/C++ Precompiler.	proc
Oracle Developer demonstration showing how to call Oracle8i Time Series functions in an Oracle Forms application.	dev2k

The `README.txt` file in the demonstration directory introduces the demonstrations. In addition, the directory for each demonstration contains a `README.txt` file with a more detailed description of that demonstration.

**See Also:** See the *Oracle8i Time Series User's Guide* for more information on Oracle Time Series demonstrations.

## Oracle8i Visual Information Retrieval

A sample program is included with Visual Information Retrieval. The sample program demonstrates how to load two images into the database, generate their signatures, and then compare their signatures using a weighted similarity function.

The sample program uses two data files, `virdemo1.dat` and `virdemo2.dat`, as its input. No other input or parameters are required.

### Environment

The following assumptions are made about the environment:

- Visual Information Retrieval is installed and PUBLIC has EXECUTE privilege on it.
- The installation script has been run. The VIRDEMODIR directory has been created and granted PUBLIC READ access so that the image data file can be read into the database.
- The `virdemo1.dat` and `virdemo2.dat` files are valid image files that are located in the VIRDEMODIR directory and the user has read and write access to the directory.
- User SCOTT has the default TIGER password. You may need to increase the tablespace allocated to SCOTT to successfully run this sample program.

## Running the Sample Program

There are two ways to run the sample program: using the included sample images, or using your own images.

[Example 6-1](#) runs the sample program using the included image files. The images are compared using equal attribute weights:

- Globalcolor = 1.0
- Localcolor = 1.0
- Texture = 1.0
- Structure = 1.0

### **Example 6-1** *Run the Sample Program with Included Images*

```
% virdemo  
Image 1 and 2 have a similarity score of 0.0
```

[Example 6-2](#) shows how to specify your own images on the command line. The images be located in the `$ORACLE_HOME/ord/vir/demo` directory.

### **Example 6-2** *Run the Sample Program with Your Own Images*

```
% virdemo image1 image2 global_color local_color texture structure
```

You must specify all six parameters; the two file names and the four attribute weights (ranging from 0.0 to 1.0) in this sample program. Note that when using the `VIRScore()` operator in your own applications, you only need to provide one attribute weight.

The Visual Information Retrieval demonstration directory provides several other sample image files to demonstrate the effects of emphasizing the different visual attributes. You can use an image viewer (such as `xv`) to display the images, and then compare them using the sample program, experimenting with different weights.

**See Also:** See the *Oracle8i Visual Information Retrieval User's Guide and Reference* and *Oracle8i Visual Information Retrieval Java Client User's Guide and Reference* for more information on Visual Information Retrieval.

## Oracle8i Spatial

For information on Oracle8i Spatial, see the `$ORACLE_HOME/md/demo/readme.txt` file and the *Oracle8i Spatial User's Guide and Reference*.



---

---

## Optimal Flexible Architecture

This appendix contains information on Optimal Flexible Architecture. It contains the following sections:

- [Optimal Flexible Architecture](#)
- [Optimal Flexible Architecture Implemented on UNIX](#)

## Optimal Flexible Architecture

Oracle Corporation recommends that the Optimal Flexible Architecture (OFA) standard be implemented when installing and configuring Oracle8i. The OFA standard is a set of configuration guidelines created to ensure fast, reliable Oracle databases that require little maintenance.

OFA is designed to:

- Organize large amounts of complicated software and data on disk to avoid device bottlenecks and poor performance
- Facilitate routine administrative tasks such as software and data backup, which are often vulnerable to data corruption
- Facilitate switching between multiple Oracle databases
- Adequately manage and administer database growth
- Help eliminate fragmentation of free space in the data dictionary, isolate other fragmentation, and minimize resource contention

## Characteristics of an OFA-Compliant Database

An OFA-compliant database provides the following benefits:

### File System Organization

The file system is organized to enable easy administration for issues such as:

- Adding data into existing databases
- Adding users
- Creating databases
- Adding hardware

### Distributed I/O Loads

I/O loads are distributed across enough disk drives to prevent performance bottlenecks.

### Hardware Support

In most cases, investment in new hardware is not required to take advantage of the OFA standard.

### **Safeguards Against Drive Failures**

By distributing applications across more than one drive, drive failures affect as few applications as possible.

### **Distribution of Home Directories**

The following items can be distributed across more than one disk drive:

- The collection of home directories
- The contents of an individual home directory

### **Integrity of Login Home Directories**

It is possible to add, move, or delete login home directories without having to revise programs that refer to them.

### **Independence of UNIX Directory Subtrees**

Categories of files are separated into independent UNIX directory subtrees so that files in one category are minimally affected by operations on files in other categories.

### **Supports Concurrent Execution of Application Software**

You can execute multiple versions of application software simultaneously, enabling the user to test and use a new release of an application before abandoning the previous version. Transferring to a new version after an upgrade is simple for the administrator and transparent for the user.

### **Distinguishes Administrative Information for Each Database**

The ability to separate administrative information on one database from that of another ensures a reasonable structure for the organization and storage of administrative data.

### **Uses Consistent Database File Naming**

Database files are named so that:

- Database files are easily distinguishable from all other files
- Files of one database are easily distinguishable from files of another database
- Control files, redo log files, and data files are identifiable as such
- The association of data file to tablespace is clearly indicated

### **Separation of Tablespace Contents**

Tablespace contents are separated to:

- Minimize tablespace free space fragmentation
- Minimize I/O request contention
- Maximize administrative flexibility

### **I/O Loads Tuned Across All Drives**

I/O loads are tuned across all drives, including drives storing Oracle data in raw devices.

## **Optimal Flexible Architecture Implemented on UNIX**

A careful naming strategy for database files eliminates data administration problems. The OFA rules provided here correspond to the original OFA recommendations published in *The OFA Standard: Oracle8 for Open Systems*.

## **Mount Points**

This section describes the naming conventions for mount points.

### **Create Mount Points**

An installation of Oracle8i requires at least two mount points: one for the software and at least one for the database files. If implementing the recommended OFA, at least four mount points are required: one for the software and at least three for database files.

## Mount Point Syntax

Name all mount points using the syntax `/pm`, where *p* is a string constant and *m* is a unique fixed-length key (typically a two-digit number) used to distinguish each mount point. For example: `/u01` and `/u02`, or `/disk01` and `/disk02`.

## Naming Mount Points for Very Large Databases (VLDBs)

If each disk drive contains database files from one application and there are enough drives for each database to ensure no I/O bottleneck, then use the syntax `/q/dm` for naming mount points. [Table A-1](#) describes the variables used in this syntax.

**Table A-1 Syntax for Naming Mount Points**

Variable	Description
<i>q</i>	A string denoting that Oracle data is stored here
<i>dm</i>	The value of the initialization parameter DB_NAME (synonymous with the instance <i>sid</i> for single-instance databases)

For example, mount points named `/u01/oradata/test` and `/u02/oradata/test` allocate two drives for the Oracle test database.

## Naming Directories

This section describes the naming conventions for OFA compliant directories.

### Home Directory Syntax

Name home directories using the syntax `/pm/h/u`. [Table A-2](#) describes the variables used in this syntax.

**Table A-2 Syntax for Naming Home Directories**

Variable	Description
<i>pm</i>	A mount point name
<i>h</i>	A standard directory name
<i>u</i>	The name of the owner of the directory

For example, `/u01/app/oracle` is the Oracle Server software owner home directory (also referred to as the Oracle base directory and the default used by the Installer) and `/u01/app/applmgr` is an Oracle applications software owner home directory.

Placing home directories at the same level in the UNIX file system is advantageous because it allows the collection of applications owner login home directories on different mount points to be referred to with the single pattern matching string, `/* /app/*`.

### Referring to Path Names

Refer to explicit path names only in files designed specifically to store them, such as `/etc/passwd` and the Oracle `oratab` file. Refer to group memberships only in the `/etc/group` file.

### Software Directories

To help fulfill the OFA feature of simultaneously executing multiple versions of application software, store each version of the Oracle8i Server software in a directory matching the pattern `/pm/h/product/v`.

[Table A-3](#) describes the variables used in this syntax.

**Table A-3 Syntax for Naming Oracle8i Server Software Directories**

Variable	Description
<i>pm</i>	A mount point name
<i>h</i>	A standard directory name
<i>v</i>	The version of the software

For example, `/u01/app/oracle/product/817` indicates the start of the directory structure where the Oracle8i Server files are located. Set the `ORACLE_HOME` environment variable to this directory.

## Naming Subdirectories

To facilitate the organization of administrative data, Oracle Corporation recommends that you store database-specific administration files in subdirectories matching the pattern *h/admin/d/a/*, where *h* is the Oracle software owner's home directory, *d* is the database name (DB\_NAME), and *a* is a subdirectory for each of the database administration files. Table A-4 describes the database administration file subdirectories:

**Table A-4 Subdirectories for Database Administration Files**

Sub Directory	Description
adhoc	Ad hoc SQL scripts for a particular database
arch	Archived redo log files
adump	Audit files (Set AUDIT_FILE_DEST in the <code>configdb_name.ora</code> file to the <code>/adump</code> directory. Clean out this subdirectory periodically).
bdump	Background process trace files
cdump	Core dump files
create	Programs used to create the database
exp	Database export files
logbook	Files recording the status and history of the database
pfile	Instance parameter files
udump	User SQL trace files

For example, the subdirectory `adhoc` would have the following path name, `/u01/app/oracle/admin/sab/adhoc/` if it were part of the database named `sab`.

## Naming Database Files

The following naming convention for database files ensures that they are easily identifiable:

- For control files, use `/pm/q/d/control.ct1`
- For redo log files, use `/pm/q/d/redon.log`
- For data files, use `/pm/q/d/tn.dbf`

[Table A-5](#) describes this syntax:

**Table A-5** Syntax for Naming Database Files

Variable	Description
<i>pm</i>	A mount point name described previously in this chapter
<i>q</i>	A string distinguishing Oracle data from all other files (usually named ORACLE or oradata)
<i>d</i>	The DB_NAME of the database
<i>t</i>	An Oracle tablespace name
<i>n</i>	A two-digit string

---

---

**Note:** Do not store files other than a control, redo log, or data file associated with database *d* in the path `/pm/q/d`.

---

---

Following this convention, you could produce, for example, a data file with the name `/u03/oradata/sab/system01.dbf`, making it easy to see to which database the file belongs.

## Separate Segments with Different Requirements

Separate groups of segments with different lifespans, I/O request demands, and backup frequencies across different tablespaces.

[Table A-6](#) describes the special tablespaces that you must create for each Oracle database. These tablespaces are in addition to those required for application segments.

**Table A-6 Special Tablespaces**

Tablespace	Description
SYSTEM	Data dictionary segments
TEMP	Temporary segments
RBS	Rollback segments
USERS	Miscellaneous user segments
INDX	Index associated with data in USERS tablespace
OEM_REPOSITORY	repository for Oracle Enterprise Manager
DRSYS	Oracle <i>interMedia</i> segment

This method is effective because dictionary segments are never dropped, and no other segments that can be dropped are allowed in the SYSTEM tablespace. This ensures that the SYSTEM tablespace does not require a rebuild due to tablespace free space fragmentation.

Because rollback segments are not stored in tablespaces holding applications data, the administrator is not blocked from taking an application's tablespace offline for maintenance. The segments are partitioned physically by type, and the administrator can record and predict data growth rates without complicated tools.

## Naming Tablespaces

Name tablespaces descriptively using a maximum of eight characters. Although Oracle8i tablespace names can be 30 characters long, portable UNIX file names are restricted to 14 characters. The recommended standard for a data file basename is *tn.dbf*, where *t* is a descriptive tablespace name and *n* is a two-digit string. Because the extension plus the two-digit string occupy a total of six characters, only eight characters remain for the tablespace name.

Descriptive names enable the data file to be associated with the tablespace that uses it. For example, the names GLD and GLX might be used for the tablespaces storing General Ledger data and indices, respectively.

---

---

**Note:** Do not embed reminders of the word "tablespace" in your tablespace names. Tablespaces are distinguishable by context, and names do not need to convey information on type.

---

---

## Exploiting the OFA Structure for Oracle Files

Table A-7 describes the syntax used for identifying classes of files.

**Table A-7** Directory Structure Syntax for Identifying Classes of Files

Directory Structure Syntax	Description
/u[0-9][0-9]	User data directories
*/home/*	User home directories
*/app/*	User application software directories
*/app/applmgr	Oracle apps software subtrees
*/app/oracle/product	Oracle Server software subtrees
*/app/oracle/product/8.1.7	Oracle Server 8.1.7 distribution files
*/app/oracle/admin/sab	sab database administrative subtrees
*/app/oracle/admin/sab/arch/*	sab database archived log files
*/oradata	Oracle data directories
*/oradata/sab/*	sab database files
*/oradata/sab/*.log	sab database redo log files

## OFA File Mapping

**Table A-8** shows a hierarchical file mapping of a sample OFA-compliant database, including each file mount point, application, database, and tablespace. The file names indicate the file type (control, log, or data).

**Table A-8 Hierarchical File Mapping for OFA Installation**

Directory	Description
/	Root mount point
/u01/	User data mount point 1
/u01/app/	Subtree for app software
/u01/app/oracle/	Home for oracle software user
/u01/app/oracle/admin/	Subtree for database admin files
/u01/app/oracle/admin/TAR/	Subtree for Support logs
/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/local/	Subtree for local Oracle software
/u01/app/oracle/local/aps6/	An Oracle6 admin package
/u01/app/oracle/local/aps7/	An Oracle7 admin package
/u01/app/oracle/product/	Distribution files
/u01/app/oracle/product/8.0.6/	ORACLE_HOME for 8.0.6 instances
/u01/app/oracle/product/8.1.6/	ORACLE_HOME for 8.1.6 instances
/u01/app/ltb/	Home directory for a user
/u01/app/sbm/	Home directory for a user
/u01/oradata/	Subtree for Oracle data
/u01/oradata/db_name1/	Subtree for <i>db_name1</i> database files
/u01/oradata/db_name2/	Subtree for <i>db_name2</i> database files
/u02/	User data mount point 2
/u02/home/	Subtree for login home directories
/u02/home/cvm/	Home directory for a user
/u02/home/vrm/	Home directory for a user

**Table A-8 Hierarchical File Mapping for OFA Installation (Cont.)**

<b>Directory</b>	<b>Description</b>
/u02/oradata/	Subtree for Oracle data
/u02/oradata/db_name1/	Subtree for <i>db_name1</i> database files
/u02/oradata/db_name2/	Subtree for <i>db_name2</i> database files
/u03/	User data mount point 3
/u03/oradata/	Subtree for Oracle data
/u03/oradata/db_name1/	Subtree for <i>db_name1</i> database files
/u03/oradata/db_name2/	Subtree for <i>db_name2</i> database files

## Raw Device Sizes

Choose a small set of standard sizes for all raw devices that can be used to store Oracle database files. In general, standardizing on a single size is recommended. If a single size is used, raw files can be moved from one partition to another safely. The size should be small enough so that a fairly large number can be created but large enough to be convenient.

For example, a 2 GB drive could be divided into 10 partitions of 200 MB each. This is a good balance between size and number. Any tablespace using raw devices should stripe them across several drives. If possible, use a logical volume manager to do the striping.

## Directory Structure

The following section describes the directory structure for OFA compliant installations.

### Oracle Base Directory

The Oracle base directory is the root of the Oracle directory structure. When installing an OFA-compliant database using the Oracle Universal Installer, the default Oracle base directory is set to `/pm/app/oracle`. [Table A-9](#) describes an Oracle base directory structure and content.

**Table A-9 Oracle Base Directory Structure and Content**

Directory Structure Syntax	Description
admin	Administrative files
doc	Online documentation
local	Subtree for local Oracle software
product	Oracle software

### Oracle Home Directory

If you install an OFA-compliant Oracle Server, the Oracle home directory is `/pm/app/oracle/product/release_number`. [Table A-10](#) describes the Oracle home directory structure and content. Under UNIX, the Oracle home directory contains the subdirectories described in [Table A-10](#), as well as a subdirectory for each Oracle product installed.

**Table A-10 Oracle Home Directory Structure and Content**

Directory	Description
assistants	Configuration Assistants
bin	Binaries for all products
ctx	<i>interMedia</i> Text options
dbs	<i>init</i> <i>sid</i> .ora, <i>lksid</i>
install	Installation related files
lib	Oracle product libraries
jlib	Java classes
md	Spatial options

**Table A–10 Oracle Home Directory Structure and Content (Cont.)**

Directory	Description
mlx	Xerox Stemmer (for <i>interMedia</i> Text options)
network	Net8
nlsrtl	NLS runtime loadable data
ocommon	Common files for all products
odg	Data gatherer
oracore	Core libraries
ord	Data options
otrace	Oracle TRACE
plsql	PL/SQL
precomp	Precompilers
rdbms	Server files and libraries required for the database
slax	SLAX parser
sqlplus	SQL*Plus

### Contents of Product Subdirectories

[Table A–11](#) describes the subdirectories contained in each product subdirectory:

**Table A–11 Contents of Product Subdirectories**

Directory	Description
admin	Administrative SQL and shell scripts (for example, <code>catalog.sql</code> , <code>catexp.sql</code> , and <code>demo.sql</code> )
admin/*	Special directories for other products
admin/resource	Resource files
admin/terminal	Runtime terminal files
demo	Demonstration scripts and datafiles
doc	README files (for example, <code>readmeunix.doc</code> )
install	Product installation scripts
jlib	Product Java classes
lib	Product libraries and distributed makefiles
log	Trace files and log files (for example, <code>orasrv.log</code> and <code>*.trc</code> files)

**Table A–11 Contents of Product Subdirectories (Cont.)**

Directory	Description
mesg	U.S. message files and Multilingual Option (formerly National Language Support) message text and binary files (for example, <code>oraus.msg</code> and <code>oraus.msb</code> )

## Examples of Product Subdirectories

[Table A–12](#) shows examples of product subdirectories and their contents.

**Table A–12 Examples of Product Subdirectories**

Directory	Description
rdbms	admin, doc, install, lib, log, mesg
sqlplus	,admin, demo, doc, install, lib, mesg

## File Naming Conventions in the admin Directory

[Table A–13](#) shows the SQL scripts located in the `rdbms/admin` directory.

**Table A–13 admin Directory, File Naming Conventions**

File	Description
cat*.sql	Creates catalog and data dictionary tables and views. The following files are run automatically during all installation: <code>catalog.sql</code> <code>catproc.sql</code> <code>catrep.sql</code> The <code>catproc.sql</code> file in turn runs the scripts for creating the standard PL/SQL packages, such as <code>DBMS_SQL</code> and <code>DBMS_OUTPUT</code> .
d*.sql	Downgrade scripts
dbms*.sql	Additional database packages
u*.sql	Upgrade scripts
utl*.sql	Creates tables and views for database utilities

## Filename Extensions

[Table A-14](#) describes filename extensions.

**Table A-14** *Filename Extensions*

Extension	Description
.a	Object file libraries; Ada runtime libraries
.aud	Oracle audit file
.bdf	X11 font description file
.bmp	X11 bitmap file
.c	C source file
.ctl	SQL*Loader control file; Oracle Server control file
.dat	SQL*Loader datafile
.dbf	Oracle Server tablespace file
.dmp	Export file
.doc	ASCII text file
.env	Shell script file for setting environment
.h	C header file; also, <code>sr.h</code> is a SQL*Report Writer help file
.jar	Java class archive
.l	UNIX manual page
.lis	Output of SQL*Plus scripts
.log	Installation log files; Oracle Server redo log files
.mk	make files
.msb	NLS message file (binary)
.msg	NLS message file (text)
.o	Object module
.ora	Oracle configuration files
.orc	Installation prototype files
.pc	Pro*C source file
.pco	Pro*COBOL source file
.ppd	printer driver file
.sh	Bourne shell script file

**Table A-14** *Filename Extensions (Cont.)*

<b>Extension</b>	<b>Description</b>
.sql	SQL* script files
.sys	Bourne shell script file
.tab	SQL* script file
.trc	Trace files
.tut	Bourne shell script file
.utd	Uniform Terminal Definitions
.zip	Zip file

---



---

---

# Index

## Symbols

---

- ? abbreviation, 1-2
- @ abbreviation, 1-2

## A

---

- accounts, 1-15
- adapters utility, 5-3
- ADDRESS specification protocols, 5-6
- administering SQL, 3-2
- advanced-developer demo, 6-8
- asynchronous I/O, 2-5
  - using, 2-5
- authorization, 1-17
- automatic login
  - remote\_os\_roles, 1-18

## B

---

- BEQ protocol, 5-6
  - syntax for connection parameters, 5-7
- binding processes, 2-8
- block size, 2-10
- buffer cache size
  - tuning, 2-11
- buffer manager, 2-3

## C

---

- C
  - Pro\*C/C++, 4-4
- cache size, 2-11
  - tuning, 2-11

- caching
  - SGA tuning, 2-14
- catching routine, 4-14
  - example, 4-14
- CATPROC.SQL, A-15
- changing databases, 1-8
- command interpreter, 1-7
- commands
  - orapwd, 1-18
  - sar, 2-4
  - vmstat, 2-4
- common environment
  - setting, 1-8
- configuration
  - encapsulator, 5-11
- configuration files
  - Net8, 5-2
  - precompiler, 4-2
- CONNECT INTERNAL, 1-17
- coraenv, 1-4
- CPU usage
  - priority level of processes, 2-8
  - processor binding, 2-8
  - single-task architecture, 2-9
  - tuning, 2-8

## D

---

- data options, 6-1
- database files, 1-17
- database writer
  - tuning, 2-5
- database writer tuning, 2-5
- dba group

- members, 1-16
- dbhome file, 1-8
- DBWR tuning, 2-5
- debugger programs, 4-3
- demonstration
  - precompiler, 1-26
  - the procedural option, PL/SQL, 1-26
- demonstration programs
  - oracle call interface, 4-9
  - Pro\*C/C++, 4-5
- demonstration tables
  - creating manually, 3-3
  - deleting, 3-4
  - SQL\*Plus, 3-3
- demos (demonstration files)
  - advanced-developer, 6-8
  - electric utility, 6-8
  - Java, 6-8
  - Oracle Call Interface (OCI), 6-9
  - Oracle Developer, 6-9
  - PRO\*C/C++, 6-8
  - retrofitting existing tables, 6-8
  - usage, 6-8
- Developer/2000 (Oracle Developer) demo, 6-9
- disk
  - monitoring performance, 2-7
  - quotas, 1-14
- disk I/O
  - asynchronous I/O, 2-5
  - file system type, 2-7
  - I/O slaves, 2-5
  - tuning, 2-5
  - tuning the database writer, 2-5
- disk performance, 2-7
  - issues, 2-7
- documentation
  - accessing, xiii
  - navigation files, xiii
  - related, xiv
  - where installed, xiii
- driver
  - post-wait, 2-8
- dynamic and static linking
  - Oracle libraries, 4-3

---

## E

- editor SQL\*Plus, 3-5
- electric utility demo, 6-8
- encapsulator configuration, 5-11
- environment variables, 1-2
  - LD\_LIBRARY64\_PATH, 1-6
  - LD\_LIBRARYN32\_PATH, 1-6
  - NLS\_LANG, 1-3
  - ORA\_NLS, 1-3
  - ORACLE\_DOC, xiii
  - ORACLE\_HOME, 1-3
  - ORACLE\_SID, 1-4
  - ORACLE\_TRACE, 1-4
  - ORAENV\_ASK, 1-4
  - SGL\_ABI, 1-6
  - SHELL, 1-7
  - SRCHOME, 1-7
  - THREADS\_FLAG, 1-7
  - TNS\_ADMIN, 1-4, 5-2
  - TWO\_TASK, 1-5
- exports, 2-9

---

## F

- file size, 2-10
- file systems, 2-7
- files
  - dbhome, 1-8
  - oratab, 1-28
  - root.sh, 1-8
  - trace files, 2-15

---

## G

- glogin.sql file, 3-2

---

## H

- help facility, 3-4

---

## I

- imports, 2-9
- initialization parameters
  - BACKGROUND\_DUMP\_DEST, 1-20, 2-15

- BITMAP\_MERGE\_AREA\_SIZE, 1-20
- COMMIT\_POINT\_STRENGTH, 1-20
- CONTROL\_FILES, 1-20
- CREATE\_BITMAP\_AREA\_SIZE, 1-20
- DB\_BLOCK\_BUFFERS, 1-20
- DB\_BLOCK\_SIZE, 1-20
- DB\_FILE\_DIRECT\_IO\_COUNT, 1-20
- DB\_FILE\_MULTIBLOCK\_READ\_COUNT, 1-20
- DB\_FILES, 1-20
- DISTRIBUTED\_TRANSACTIONS, 1-20
- HASH\_AREA\_SIZE, 1-20
- HASH\_MULTIBLOCK\_IO\_COUNT, 1-20
- IRIX\_CPU\_AFFINITY, 2-12
- IRIX\_SCHEDULER, 2-13
- JAVA\_POOL\_SIZE, 1-20
- LOCK\_SGA, 1-20
- LOG\_ARCHIVE\_DEST, 1-20
- LOG\_ARCHIVE\_FORMAT, 1-20
- LOG\_BUFFER, 1-20
- LOG\_CHECKPOINT\_INTERVAL, 1-20
- MTS\_LISTENER\_ADDRESS, 1-20
- MTS\_MAX\_DISPATCHERS, 1-20
- MTS\_MAX\_SERVERS, 1-20
- MTS\_SERVERS, 1-20
- NLS\_LANGUAGE, 1-21
- NLS\_TERRITORY, 1-21
- OBJECT\_CACHE\_MAX\_SIZE\_PERCENT, 1-21
- OBJECT\_CACHE\_OPTIMAL\_SIZE, 1-21
- OPEN\_CURSORS, 1-21
- OS\_AUTHENT\_PREFIX, 1-21
- PROCESSES, 1-21
- SHARED\_POOL\_SIZE, 1-21
- SORT\_AREA\_SIZE, 1-21
- initsid.ora file, 1-19
- Intelligent Agent
  - configuration on SNMP, 5-11
- interMedia
  - Java demo, 6-6
  - MediaAnnotator, 6-6
  - sample audio scripts, 6-2
  - sample programs, 6-3
  - sample video scripts, 6-5
  - Web Agent and Clipboard, 6-8
- interMedia Audio, Video, and Image, 6-2
- interMedia Locator, 6-6

- interMedia Text, 6-2
- interrupting SQL\*Plus, 3-6
- I/O
  - tuning, 2-5
- I/O slaves, 2-5
- IPC protocol, 5-8
- ireclen, 4-3
- IRIX\_CPU\_AFFINITY parameter, 2-12
- IRIX\_SCHEDULER parameter, 2-13

## J

---

- Java demos, 6-8

## L

---

- language, 1-3
- LD\_LIBRARY64\_PATH variable, 1-6
- LD\_LIBRARYN32\_PATH variable, 1-6
- limits
  - resource, 1-14
- linking, 2-9
- log writer process (LGWR), 2-8
- login.sql file, 3-2

## M

---

- memory
  - estimating usage, 1-13
  - SGA tuning, 2-5
  - tuning, 2-3
  - virtual, 1-13
- memory management, 2-3
  - control paging, 2-4
  - single shared memory segment, 2-5
  - swap space, 2-4
  - UNIX kernel, 2-5
- multiple signal handlers, 4-13
- multi-thread server, 5-4

## N

---

- National Language Support (NLS) variable, 1-3
- Net8
  - adapters utility, 5-3
  - ADDRESS specification, 5-5

- BEQ protocol, 5-6
- files and utilities, 5-2
- IPC protocol, 5-8
- multi-thread server, 5-4
- oracle advanced security, 5-13
- oracle connection manager, 5-3
- products and features, 5-2
- protocol support, 5-5
- protocols, 5-5
- RAW protocol, 5-9
- TCP/IP protocol, 5-10
- Net8 configuration files, 5-2
- NLS\_LANG variable, 1-3

## O

---

- OCI demo, 6-9
- OFA
  - See also Optimal Flexible Architecture
- Optimal Flexible Architecture, A-1
  - characteristics of OFA, A-2
  - file mapping, A-2, A-11
  - implementation on Oracle8 for Unix, A-2
- ORA\_NLS variable, 1-3
- Oracle Advanced Security
  - .bak files, 5-13
  - security and single sign-on, 5-13
- oracle advanced security, 5-13
  - .bak files, 5-13
  - security and single sign-on, 5-13
- oracle call interface, 4-9
  - demonstration programs, 4-9
- Oracle Call Interface (OCI) demo, 6-9
- oracle connection manager, 5-3
- Oracle Data Option Demos, 6-1
- Oracle Developer demo, 6-9
- Oracle environment variables
  - EPC\_DISABLED, 1-3
  - NLS\_LANG, 1-3
  - ORA\_NLS33, 1-3
  - ORACLE\_BASE, 1-3
  - ORACLE\_HOME, 1-3
  - ORACLE\_PATH, 1-4
  - ORACLE\_SID, 1-4
  - ORACLE\_TRACE, 1-4

- ORAENV\_ASK, 1-4
- TNS\_ADMIN, 1-4
- TWO\_TASK, 1-5
- UNIX variables, 1-5
- Oracle Forms demo, 6-9
- Oracle HTTP Server, 1-23
- Oracle Information Navigator, xiii
- Oracle libraries
  - static and dynamic linking, 4-3
- oracle precompiler and OCI linking and makefiles
  - custom makefiles, 4-11
  - undefined symbols, 4-11
- Oracle Server accounts, 1-15
- oracle software owner, 1-15
  - special accounts, 1-15
- Oracle System Identifier, 1-4
- Oracle TCP/IP protocol, 5-10
- ORACLE\_DOC variable, xiii
- ORACLE\_HOME environment variable, 1-3
- ORACLE\_HOME variable, 1-2
- ORACLE\_SID
  - environment variable, 1-4
- ORACLE\_SID variable, 1-2
- Oracle8i Spatial, 6-11
- oraenv file
  - description, 1-8
  - moving between databases, 1-8
- ORAENV\_ASK, 1-4
- orapwd command, 1-18
- oratab file, 1-28
- oreclen, 4-3

## P

---

- page-out activity, 2-4
- paging, 2-4
- paging space, 2-3
  - tuning, 2-3, 2-4
- passwords, 1-18
- PL/SQL demonstrations
  - loading, 1-26
- post-wait driver
  - tuning, 2-8
- precompiler
  - running demonstration, 1-26

- value of ireclen and oreclen, 4-3
- precompiler configuration files, 4-2
- precompilers
  - overview, 4-2
  - signals, 4-13
  - uppercase to lowercase conversion, 4-3
  - values, 4-3
  - vendor debugger programs, 4-3
- Pro\*C/C++
  - demonstration programs, 4-5
  - makefiles, 4-5
  - signals, 4-13
  - using, 4-4
- PRO\*C/C++ demo, 6-8
- processor affinity/binding, 2-8
- PRODUCT\_USER\_PROFILE Table, 3-3
- protocols, 5-5
  - ADDRESS specification, 5-6

## R

---

- raw devices
  - buffer cache size, 2-11
  - disadvantages, 2-16
  - guidelines for using, 2-17
  - raw disk partition availability, 2-17
  - setting up, 2-17
- RAW protocol, 5-9
- README.txt files, 6-9
- related documentation, xiv
- relinking, 1-9
- remote connections parameters
  - OS\_AUTHENT\_PREFIX, 1-18
  - REMOTE\_OS\_AUTHENT, 1-18
  - REMOTE\_OS\_ROLES, 1-18
- remote passwords, 1-18
- resource contention
  - tuning, 2-9
- resource limits, 1-14
- restrictions (SQL\*Plus), 3-7
  - resizing windows, 3-7
  - return codes, 3-7
- retrofit demo, 6-8
- root user, 1-15
- root.sh script, 1-8

- running operating system commands, 3-6

## S

---

- sar, 2-2
- sar command, 2-4, 2-11
- scripts
  - start\_peer, 5-12
- security, 1-16, 1-17
  - CONNECT INTERNAL, 1-17
  - features of UNIX, 1-16
  - file ownership, 1-16
  - group accounts, 1-16
  - Server Manager access, 1-17
  - SHUTDOWN command, 1-17
  - STARTUP command, 1-17
  - two-task architecture, 1-16
- semaphores
  - replacing with post-wait driver, 2-8
- Server Manager
  - commands, 1-17
  - security, 1-17
- setting up raw devices, 2-17
- setup files
  - SQL\*Plus, 3-2
- SGA
  - calculating, 1-11
  - locking in physical memory, 2-14
  - relocating, 1-11, 1-12
  - tuning, 2-5
- SGL\_ABI variable, 1-6
- shadow process, 1-16
- shared memory
  - SGA tuning, 2-5
- SHELL, 1-7
- SHUTDOWN command, 1-17
  - security, 1-17
- SIGCLD two-task signal, 4-13
- SIGINT two-task signal, 4-13
- SIGIO two-task signal, 4-13
- signal handlers
  - using, 4-12
- signal routine, 4-14
  - example, 4-14
- signals

- two-task, 4-13
- SIGPIPE two-task signal, 4-13
- SIGTERM two-task signal, 4-13
- SIGURG two-task signal, 4-13
- Simple Network Management Protocol (SNMP), 5-11
- single shared memory segment, 2-5
- site profile SQL\*Plus, 3-2
- SNMP
  - and Intelligent Agent, 5-11
- SNMP and Intelligent Agent, 5-11
- snmpd executable, 5-12
- software distribution, 1-3
- Spatial demonstrations, 6-11
- special groups
  - root, 1-16
- SPOOL command
  - SQL\*Plus, 3-7
  - using, 3-7
- SQL scripts, 2-3
- SQL\*Plus
  - default editor, 3-5
  - demonstration tables, 3-3
  - editor, 3-5
  - help facility, 3-4
  - interrupting, 3-6
  - PRODUCT\_USER\_PROFILE Table, 3-3
  - restrictions, 3-7
  - running operating system commands, 3-6
  - setup files, 3-2
  - site profile, 3-2
  - SPOOL command, 3-7
  - system editor, 3-5
  - user profile, 3-2
  - using, 3-5
- SRCHOME variable, 1-7
- start\_peer script, 5-12
- STARTUP command, 1-17
- static and dynamic linking
  - oracle libraries, 4-3
- superuser, 1-15
- swap space, 2-3
  - tuning, 2-3
- SYS account, 1-15
  - privileges, 1-15

- SYSDATE, 1-9
- SYSTEM account, 1-15
  - privileges, 1-15
- system editor
  - SQL\*Plus, 3-5
- System Global Area or Shared Global Area (SGA)
  - relocating, 1-11, 1-12
- system time, 1-9

## T

---

- TCP/IP protocol, 5-10
  - ADDRESS, 5-10
- thread support, 4-12
- THREADS\_FLAG variable, 1-7
- TNS listener, 5-10
  - configuring for Oracle TCP/IP protocol, 5-10
- TNS\_ADMIN environment variable, 1-4
- tools, 2-2
  - sar, 2-2
- trace and alert files
  - alert files, 2-15
  - trace file names, 2-15
  - using, 2-11, 2-15
- tracing Bourne shell scripts, 1-4
- tuning, 2-3, 2-5, 2-9
  - block and file size, 2-10
  - buffer cache size, 2-11
  - client/server configuration, 2-8
  - CPU usage, 2-8
  - disk I/O, 2-5
  - I/O bottlenecks, 2-5
  - memory management, 2-3
  - post-wait driver, 2-8
  - resource contention, 2-9
  - trace and alert files, 2-11, 2-15
- TWO\_TASK environment variable, 1-5
- two-task, 1-16
  - signals, 4-13

## U

---

- ufs file system type, 2-7
- UNIX Environment Variables Used with Oracle8i, 1-3

usage demo, 6-8  
user interrupt handler, 4-14  
user profile  
    SQL\*Plus, 3-2  
using SQL\*Plus, 3-5  
utility (electric) demo, 6-8

## **V**

---

vmstat command, 2-4

## **W**

---

writer activity, 2-5

## **X**

---

XA functionality, 4-15

