

# Oracle8i

SQL Reference

Release 2 (8.1.6)

December 1999

Part No. A76989-01

**ORACLE**

---

SQL Reference, Release 2 (8.1.6)

Part No. A76989-01

Copyright © 1996, 1999, Oracle Corporation. All rights reserved.

Primary Author: Diana Lorentz

Contributors: Alan Downing, Alex Tsukerman, Alok Pareek, Amit Ganesh, Andrew Witkowski, Angela Amor, Ann Rhee, Aravind Yalamanchi, Ari Mozes, Arvind Nithrakashyap, Ashok Joshi, Ashwini Surpur, Bill Lee, Brian Wright, Carolyn Gray, Cetin Ozbutun, Chon Lei, Daniel Wong, Debbie Steiner, Edward Waugh, Eugene Chong, Franco Putzolu, Fred Zemke, Guhan Viswanathan, Hakkan Jakobsson, Harry Sun, Harvey Eneman, Ira Greenberg, Jack Raitto, Jags Srinivasan, Jerry Schwarz, Jianping Yang, Jim Finnerty, John Haydu, Juan Tellez, Karuna Muthiah, Lance Ashdown, Lilian Hobbs, Lois Price, Mark Johnson, Miclelle Cyran, Mohamed Zait, Mohamed Ziauddin, Muralidhar Krishnaprasad, Namit Jain, Nipun Agarwal, Peter Robson, Qin Yu, Rafi Ahmed, Randy Urbano, Ravi Murthy, Reema Al-Shaikh, Rick Anderson, Robert Jenkins, Rosanne Park, Ruth Baylis, Sanjay Kaluskar, Sankar Subramanian, Sashikanth Chandrasekaran, Sheri Blaisedell, Sriram Samu, Steve Harris, Steve Vivian, Subhransu Basu, Subramanian Muralidhar, Sukhjit Singh, Susan Kotsovolos, Tak Wang, Thong Bui, Thuvan Hoang, Vikas Arora, Vinay Srihari, Vishu Krishnamurthy, Vishy Karra, Wayne Smith, Wei Wang

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.

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.

Oracle is a registered trademark. All other company or product names mentioned are used for identification purposes only and may be trademarks of their respective owners.

---

---

# Contents

<b>Send Us Your Comments .....</b>	<b>xiii</b>
<b>Preface.....</b>	<b>xv</b>
<b>1 Introduction</b>	
<b>Lexical Conventions .....</b>	<b>1-4</b>
<b>2 Basic Elements of Oracle SQL</b>	
<b>Literals .....</b>	<b>2-2</b>
<b>Datatypes .....</b>	<b>2-9</b>
<b>Format Models .....</b>	<b>2-38</b>
<b>Nulls .....</b>	<b>2-54</b>
<b>Pseudocolumns .....</b>	<b>2-56</b>
<b>Comments .....</b>	<b>2-62</b>
<b>Database Objects .....</b>	<b>2-68</b>
<b>Schema Object Names and Qualifiers .....</b>	<b>2-72</b>
<b>Referring to Schema Objects and Parts .....</b>	<b>2-76</b>
<b>3 Operators</b>	
<b>Unary and Binary Operators .....</b>	<b>3-1</b>
<b>Precedence .....</b>	<b>3-2</b>
<b>Arithmetic Operators .....</b>	<b>3-3</b>
<b>Concatenation Operator .....</b>	<b>3-3</b>
<b>Comparison Operators .....</b>	<b>3-5</b>

<b>Logical Operators</b> .....	3-11
<b>Set Operators</b> .....	3-12
<b>Other Built-In Operators</b> .....	3-16
<b>User-Defined Operators</b> .....	3-16

## 4 Functions

<b>SQL Functions</b> .....	4-2
<b>ABS</b> .....	4-12
<b>ACOS</b> .....	4-13
<b>ADD_MONTHS</b> .....	4-13
<b>ASCII</b> .....	4-14
<b>ASIN</b> .....	4-15
<b>ATAN</b> .....	4-15
<b>ATAN2</b> .....	4-16
<b>AVG</b> .....	4-16
<b>BFILENAME</b> .....	4-17
<b>CEIL</b> .....	4-18
<b>CHARTOROWID</b> .....	4-19
<b>CHR</b> .....	4-19
<b>CONCAT</b> .....	4-20
<b>CONVERT</b> .....	4-21
<b>CORR</b> .....	4-22
<b>COS</b> .....	4-23
<b>COSH</b> .....	4-24
<b>COUNT</b> .....	4-24
<b>COVAR_POP</b> .....	4-26
<b>COVAR_SAMP</b> .....	4-27
<b>CUME_DIST</b> .....	4-28
<b>DENSE_RANK</b> .....	4-29
<b>DEREF</b> .....	4-30
<b>DUMP</b> .....	4-31
<b>EMPTY_[B   C]LOB</b> .....	4-33
<b>EXP</b> .....	4-33
<b>FIRST_VALUE</b> .....	4-34
<b>FLOOR</b> .....	4-35

<b>GREATEST</b> .....	4-36
<b>GROUPING</b> .....	4-36
<b>HEXTORAW</b> .....	4-37
<b>INITCAP</b> .....	4-38
<b>INSTR</b> .....	4-38
<b>INSTRB</b> .....	4-39
<b>LAG</b> .....	4-40
<b>LAST_DAY</b> .....	4-41
<b>LAST_VALUE</b> .....	4-42
<b>LEAD</b> .....	4-44
<b>LEAST</b> .....	4-45
<b>LENGTH</b> .....	4-45
<b>LENGTHB</b> .....	4-46
<b>LN</b> .....	4-46
<b>LOG</b> .....	4-47
<b>LOWER</b> .....	4-47
<b>LPAD</b> .....	4-48
<b>LTRIM</b> .....	4-48
<b>MAKE_REF</b> .....	4-49
<b>MAX</b> .....	4-50
<b>MIN</b> .....	4-52
<b>MOD</b> .....	4-53
<b>MONTHS_BETWEEN</b> .....	4-54
<b>NEW_TIME</b> .....	4-55
<b>NEXT_DAY</b> .....	4-55
<b>NLS_CHARSET_DECL_LEN</b> .....	4-56
<b>NLS_CHARSET_ID</b> .....	4-56
<b>NLS_CHARSET_NAME</b> .....	4-57
<b>NLS_INITCAP</b> .....	4-58
<b>NLS_LOWER</b> .....	4-59
<b>NLSSORT</b> .....	4-59
<b>NLS_UPPER</b> .....	4-60
<b>NTILE</b> .....	4-61
<b>NUMTODSINTERVAL</b> .....	4-62
<b>NUMTOYMINTERVAL</b> .....	4-63

NVL .....	4-64
NVL2 .....	4-65
PERCENT_RANK .....	4-66
POWER .....	4-67
RANK .....	4-67
RATIO_TO_REPORT .....	4-68
RAWTOHEX .....	4-69
REF .....	4-70
REFTOHEX .....	4-70
REGR_ (linear regression) functions .....	4-71
REPLACE .....	4-78
ROUND (Number Function) .....	4-79
ROUND (Date Function) .....	4-79
ROW_NUMBER .....	4-80
ROWIDTOCHAR .....	4-81
RPAD .....	4-82
RTRIM .....	4-82
SIGN .....	4-83
SIN .....	4-83
SINH .....	4-84
SOUNDEX .....	4-84
SQRT .....	4-85
STDDEV .....	4-86
STDDEV_POP .....	4-87
STDDEV_SAMP .....	4-88
SUBSTR .....	4-90
SUBSTRB .....	4-91
SUM .....	4-91
SYS_CONTEXT .....	4-93
SYS_GUID .....	4-97
SYSDATE .....	4-98
TAN .....	4-98
TANH .....	4-99
TO_CHAR (date conversion) .....	4-99
TO_CHAR (number conversion) .....	4-100

TO_DATE .....	4-101
TO_LOB .....	4-102
TO_MULTI_BYTE .....	4-103
TO_NUMBER .....	4-103
TO_SINGLE_BYTE .....	4-104
TRANSLATE .....	4-105
TRANSLATE ... USING .....	4-106
TRIM .....	4-107
TRUNC (Number Function) .....	4-108
TRUNC (Date Function) .....	4-109
UID .....	4-109
UPPER.....	4-110
USER .....	4-110
USERENV .....	4-111
VALUE .....	4-112
VAR_POP .....	4-113
VAR_SAMP .....	4-114
VARIANCE .....	4-115
VSIZE .....	4-116
ROUND and TRUNC Date Functions.....	4-117
User-Defined Functions .....	4-118

## 5 Expressions, Conditions, and Queries

Expressions .....	5-1
Conditions .....	5-14
Queries and Subqueries .....	5-19

## 6 About SQL Statements

Summary of SQL Statements .....	6-1
Finding the Right SQL Statement .....	6-5

## 7 SQL Statements

ALTER CLUSTER .....	7-2
ALTER DATABASE .....	7-6

<b>ALTER DIMENSION</b> .....	7-25
<b>ALTER FUNCTION</b> .....	7-28
<b>ALTER INDEX</b> .....	7-30
<b>ALTER JAVA</b> .....	7-45
<b>ALTER MATERIALIZED VIEW / SNAPSHOT</b> .....	7-47
<b>ALTER MATERIALIZED VIEW LOG / SNAPSHOT LOG</b> .....	7-58
<b>ALTER OUTLINE</b> .....	7-63
<b>ALTER PACKAGE</b> .....	7-64
<b>ALTER PROCEDURE</b> .....	7-67
<b>ALTER PROFILE</b> .....	7-69
<b>ALTER RESOURCE COST</b> .....	7-73
<b>ALTER ROLE</b> .....	7-76
<b>ALTER ROLLBACK SEGMENT</b> .....	7-78
<b>ALTER SEQUENCE</b> .....	7-81
<b>ALTER SESSION</b> .....	7-83
<b>ALTER SNAPSHOT</b> .....	7-100
<b>ALTER SNAPSHOT LOG</b> .....	7-101
<b>ALTER SYSTEM</b> .....	7-102
<b>ALTER TABLE</b> .....	7-123
<b>ALTER TABLESPACE</b> .....	7-176
<b>ALTER TRIGGER</b> .....	7-183
<b>ALTER TYPE</b> .....	7-186
<b>ALTER USER</b> .....	7-193
<b>ALTER VIEW</b> .....	7-198
<b>ANALYZE</b> .....	7-200
<b>ASSOCIATE STATISTICS</b> .....	7-210
<b>AUDIT <i>sql_statements</i></b> .....	7-213
<b>AUDIT <i>schema_objects</i></b> .....	7-221
<b>CALL</b> .....	7-226
<b>COMMENT</b> .....	7-228
<b>COMMIT</b> .....	7-230
<i>constraint_clause</i> .....	7-233
<b>CREATE CLUSTER</b> .....	7-254
<b>CREATE CONTEXT</b> .....	7-261
<b>CREATE CONTROLFILE</b> .....	7-263



<b>CREATE DATABASE</b> .....	7-267
<b>CREATE DATABASE LINK</b> .....	7-273
<b>CREATE DIMENSION</b> .....	7-277
<b>CREATE DIRECTORY</b> .....	7-282
<b>CREATE FUNCTION</b> .....	7-284
<b>CREATE INDEX</b> .....	7-291
<b>CREATE INDEXTYPE</b> .....	7-309
<b>CREATE JAVA</b> .....	7-311
<b>CREATE LIBRARY</b> .....	7-316
<b>CREATE MATERIALIZED VIEW / SNAPSHOT</b> .....	7-318
<b>CREATE MATERIALIZED VIEW LOG / SNAPSHOT LOG</b> .....	7-333
<b>CREATE OPERATOR</b> .....	7-339
<b>CREATE OUTLINE</b> .....	7-342
<b>CREATE PACKAGE</b> .....	7-344
<b>CREATE PACKAGE BODY</b> .....	7-348
<b>CREATE PROCEDURE</b> .....	7-353
<b>CREATE PROFILE</b> .....	7-359
<b>CREATE ROLE</b> .....	7-365
<b>CREATE ROLLBACK SEGMENT</b> .....	7-367
<b>CREATE SCHEMA</b> .....	7-369
<b>CREATE SEQUENCE</b> .....	7-371
<b>CREATE SNAPSHOT</b> .....	7-375
<b>CREATE SNAPSHOT LOG</b> .....	7-376
<b>CREATE SYNONYM</b> .....	7-377
<b>CREATE TABLE</b> .....	7-381
<b>CREATE TABLESPACE</b> .....	7-419
<b>CREATE TEMPORARY TABLESPACE</b> .....	7-424
<b>CREATE TRIGGER</b> .....	7-426
<b>CREATE TYPE</b> .....	7-437
<b>CREATE TYPE BODY</b> .....	7-447
<b>CREATE USER</b> .....	7-451
<b>CREATE VIEW</b> .....	7-456
<b>DELETE</b> .....	7-464
<b>DISASSOCIATE STATISTICS</b> .....	7-470
<b>DROP CLUSTER</b> .....	7-472

<b>DROP CONTEXT</b> .....	7-474
<b>DROP DATABASE LINK</b> .....	7-475
<b>DROP DIMENSION</b> .....	7-476
<b>DROP DIRECTORY</b> .....	7-477
<b>DROP FUNCTION</b> .....	7-478
<b>DROP INDEX</b> .....	7-480
<b>DROP INDEXTYPE</b> .....	7-482
<b>DROP JAVA</b> .....	7-483
<b>DROP LIBRARY</b> .....	7-484
<b>DROP MATERIALIZED VIEW / SNAPSHOT</b> .....	7-485
<b>DROP MATERIALIZED VIEW LOG / SNAPSHOT LOG</b> .....	7-487
<b>DROP OPERATOR</b> .....	7-489
<b>DROP OUTLINE</b> .....	7-490
<b>DROP PACKAGE</b> .....	7-491
<b>DROP PROCEDURE</b> .....	7-493
<b>DROP PROFILE</b> .....	7-494
<b>DROP ROLE</b> .....	7-495
<b>DROP ROLLBACK SEGMENT</b> .....	7-496
<b>DROP SEQUENCE</b> .....	7-497
<b>DROP SNAPSHOT</b> .....	7-498
<b>DROP SNAPSHOT LOG</b> .....	7-499
<b>DROP SYNONYM</b> .....	7-500
<b>DROP TABLE</b> .....	7-501
<b>DROP TABLESPACE</b> .....	7-503
<b>DROP TRIGGER</b> .....	7-505
<b>DROP TYPE</b> .....	7-506
<b>DROP TYPE BODY</b> .....	7-508
<b>DROP USER</b> .....	7-509
<b>DROP VIEW</b> .....	7-511
<b>EXPLAIN PLAN</b> .....	7-512
<i>filespec</i> .....	7-516
<b>GRANT <i>system_privileges_and_roles</i></b> .....	7-519
<b>GRANT <i>object_privileges</i></b> .....	7-532
<b>INSERT</b> .....	7-539
<b>LOCK TABLE</b> .....	7-547

NOAUDIT <i>sql_statements</i> .....	7-550
NOAUDIT <i>schema_objects</i> .....	7-552
RENAME .....	7-554
REVOKE <i>system_privileges_and_roles</i> .....	7-556
REVOKE <i>schema_object_privileges</i> .....	7-559
ROLLBACK .....	7-564
SAVEPOINT .....	7-567
SELECT and Subqueries .....	7-569
SET CONSTRAINT(S) .....	7-598
SET ROLE .....	7-600
SET TRANSACTION .....	7-602
<i>storage_clause</i> .....	7-605
TRUNCATE .....	7-611
UPDATE .....	7-615

## A Syntax Diagrams

## B Oracle and Standard SQL

Conformance with Standard SQL .....	B - 1
Oracle Extensions to Standard SQL .....	B - 5

## C Oracle Reserved Words



---

---

# Send Us Your Comments

## SQL Reference, Release 2 (8.1.6)

Part No. A76989-01

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:

- E-mail - [infodev@us.oracle.com](mailto:infodev@us.oracle.com)
- FAX - (650) 506-7228. Attn: Information Development
- Postal service:  
Oracle Corporation  
Server Technologies Documentation Manager  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

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

---

---

---

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



---

# Preface

*Well begun is half done.*

Aristotle, *Nicomachean Ethics*

This reference contains a complete description of the Structured Query Language (SQL) used to manage information in an Oracle database. Oracle SQL is a superset of the American National Standards Institute (ANSI) and the International Standards Organization (ISO) SQL92 standard at entry level conformance.

**See Also:**

- *PL/SQL User's Guide and Reference* for information on PL/SQL, Oracle's procedural language extension to SQL.
- *Pro\*C/C++ Precompiler Programmer's Guide*, *SQL\*Module for Ada Programmer's Guide*, and the *Pro\*COBOL Precompiler Programmer's Guide* for detailed descriptions of Oracle embedded SQL.

## Features and Functionality

*Oracle8i SQL Reference* contains information about the features and functionality of the Oracle8i and the Oracle8i Enterprise Edition products. Oracle8i and Oracle8i Enterprise Edition have the same basic features. However, several advanced features are available only with the Enterprise Edition, and some of these are optional.

**See Also:** *Getting to Know Oracle8i* for information about the differences between Oracle8i and the Oracle8i Enterprise Edition and the available features and options. That book also describes all the features that are new in Oracle8i.

## Audience

This reference is intended for all users of Oracle SQL.

## What's New in Oracle8i?

Oracle8i contains many new features, which are documented throughout this reference. For a description of all features new to this release, see *Getting to Know Oracle8i*.

### Release 8.1.5

The following top-level SQL statements are new to Release 8.1.5:

- [ALTER DIMENSION](#) on page 7-25
- [ALTER JAVA](#) on page 7-45
- [ALTER OUTLINE](#) on page 7-63
- [ASSOCIATE STATISTICS](#) on page 7-210
- [CALL](#) on page 7-226
- [CREATE CONTEXT](#) on page 7-261
- [CREATE DIMENSION](#) on page 7-277
- [CREATE INDEXTYPE](#) on page 7-309
- [CREATE JAVA](#) on page 7-311
- [CREATE OPERATOR](#) on page 7-339
- [CREATE OUTLINE](#) on page 7-342
- [CREATE TEMPORARY TABLESPACE](#) on page 4-424
- [DISASSOCIATE STATISTICS](#) on page 7-470
- [DROP CONTEXT](#) on page 7-474
- [DROP DIMENSION](#) on page 7-476
- [DROP INDEXTYPE](#) on page 7-482
- [DROP JAVA](#) on page 7-483
- [DROP OPERATOR](#) on page 7-489
- [DROP OUTLINE](#) on page 7-490



## Release 2 (8.1.6)

The following SQL functions are new to this release:

- [CORR](#) on page 4-22
- [COVAR\\_POP](#)
- [COVAR\\_SAMP](#)
- [CUME\\_DIST](#)
- [DENSE\\_RANK](#)
- [FIRST\\_VALUE](#)
- [LAG](#)
- [LAST\\_VALUE](#)
- [LEAD](#)
- [NTILE](#)
- [NUMTOYMINTERVAL](#)
- [NUMTODSINTERVAL](#)
- [PERCENT\\_RANK](#)
- [RATIO\\_TO\\_REPORT](#)
- [REGR\\_](#) (linear regression) functions
- [STDDEV\\_POP](#)
- [STDDEV\\_SAMP](#)
- [VAR\\_POP](#)
- [VAR\\_SAMP](#)

In addition, the following features have been enhanced:

- The aggregate functions have expanded functionality. See [Aggregate Functions](#) on page 4-6.
- When specifying LOB storage parameters, you can now specify caching of LOBs for read-only purposes. See "[CREATE TABLE](#)" on page 4-381.
- The section on Expressions now contains a new expression. See "[CASE Expressions](#)" on page 5-13.

- Subqueries can now be unnested. See "[Unnesting of Nested Subqueries](#)" on page 5-25.

## How this Reference Is Organized

This reference is divided into the following parts:

### Volume 1

#### **Chapter 1, "Introduction"**

This chapter defines SQL and describes its history as well as the advantages of using it to access relational databases.

#### **Chapter 2, "Basic Elements of Oracle SQL"**

This chapter describes the basic building blocks of an Oracle database and of Oracle SQL.

#### **Chapter 3, "Operators"**

This chapter describes how to use SQL operators to combine data into expressions and conditions.

#### **Chapter 4, "Functions"**

This chapter describes how to use SQL functions to combine data into expressions and conditions.

#### **Chapter 5, "Expressions, Conditions, and Queries"**

This chapter describes SQL expressions and conditions and discusses the various ways of extracting information from your database through queries.

#### **Chapter 6, "About SQL Statements"**

This chapter lists the various types of SQL statements, and provides a table to help you find the appropriate SQL statement for your database task.

### Volume 2

#### **Chapter 7, "SQL Statements"**

This chapter lists and describes all Oracle SQL statements in alphabetical order.

## Appendix A, "Syntax Diagrams"

This appendix describes how to read the syntax diagrams in this reference.

## Appendix B, "Oracle and Standard SQL"

This appendix describes Oracle compliance with ANSI and ISO standards.

## Appendix C, "Oracle Reserved Words"

This appendix lists words that are reserved for internal use by Oracle.

## Structural Changes in the Reference in Release 8.1.5

Users familiar with the Release 8.0 documentation will find that the following sections have been moved or renamed:

- The section "[Format Models](#)" now appears in Chapter 2 on page 2-38.
- Chapter 3 has been divided into several smaller chapters:
  - [Chapter 3, "Operators"](#)
  - [Chapter 4, "Functions"](#)
  - [Chapter 5, "Expressions, Conditions, and Queries"](#). The last section, "[Queries and Subqueries](#)" on page 5-19, provides background for the syntactic and semantic information in "[SELECT and Subqueries](#)" on page 7-569.
- A new chapter, [Chapter 6, "About SQL Statements"](#), has been added to help you find the correct SQL statement for a particular task.
- The *archive\_log\_clause* is no longer a separate section, but has been incorporated into "[ALTER SYSTEM](#)" on page 7-102.
- The *deallocate\_unused\_clause* is no longer a separate section, but has been incorporated into "[ALTER TABLE](#)" on page 7-123, "[ALTER CLUSTER](#)" on page 7-2, and "[ALTER INDEX](#)" on page 7-30.
- The *disable\_clause* is no longer a separate section, but has been incorporated into "[CREATE TABLE](#)" on page 4-381 and "[ALTER TABLE](#)" on page 7-123.
- The *drop\_clause* is no longer a separate section. It has become the *drop\_constraint\_clause* of the ALTER TABLE statement (to distinguish it from the new *drop\_column\_clause* of that statement). See "[ALTER TABLE](#)" on page 7-123.
- The *enable\_clause* is no longer a separate section, but has been incorporated into "[CREATE TABLE](#)" on page 4-381 and "[ALTER TABLE](#)" on page 7-123.

- The *parallel\_clause* is no longer a separate section. The clause has been simplified, and has been incorporated into the various statements where it is relevant.
- The *recover\_clause* is no longer a separate section. Recovery functionality has been enhanced, and because it is always implemented through the ALTER DATABASE statement, it has been incorporated into that section. See "[ALTER DATABASE](#)" on page 7-6.
- The sections on **snapshots** and **snapshot logs** have been moved and renamed. Snapshot functionality has been greatly enhanced, and these objects are now called **materialized views**. See "[CREATE MATERIALIZED VIEW / SNAPSHOT](#)" on page 318, "[ALTER MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-47, "[DROP MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-485, "[CREATE MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-333, "[ALTER MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-58, and "[DROP MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-487.
- The section on **subqueries** has now been combined with the SELECT statement. See "[SELECT and Subqueries](#)" on page 7-569.

## Conventions Used in this Reference

This section explains the conventions used in this book including:

- [Text](#)
- [Syntax Diagrams and Notation](#)
- [Code Examples](#)
- [Example Data](#)


### Text

The text in this reference adheres to the following conventions:

UPPERCASE	Uppercase text calls attention to SQL keywords, filenames, and initialization parameters.
<i>italics</i>	Italicized text calls attention to parameters of SQL statements.
<b>boldface</b>	Boldface text calls attention to definitions of terms.

## Syntax Diagrams and Notation

**Syntax Diagrams** This reference uses syntax diagrams to show SQL statements in [Chapter 7, “SQL Statements”](#), and to show other elements of the SQL language in [Chapter 2, “Basic Elements of Oracle SQL”](#), [Chapter 3, “Operators”](#), [Chapter 4, “Functions”](#), and [Chapter 5, “Expressions, Conditions, and Queries”](#). These syntax diagrams use lines and arrows to show syntactic structure, as shown here:



If you are not familiar with this type of syntax diagram, refer to [Appendix A, “Syntax Diagrams”](#), for a description of how to read them. This section describes the components of syntax diagrams and gives examples of how to write SQL statements. Syntax diagrams are made up of these items:

**Keywords** Keywords have special meanings in the SQL language. In the syntax diagrams, keywords appear in UPPERCASE. You must use keywords in your SQL statements exactly as they appear in the syntax diagram, except that they can be either uppercase or lowercase. For example, you must use the `CREATE` keyword to begin your `CREATE TABLE` statements just as it appears in the `CREATE TABLE` syntax diagram.

**Parameters** Parameters act as placeholders in syntax diagrams. They appear in lowercase. Parameters are usually names of database objects, Oracle datatype names, or expressions. When you see a parameter in a syntax diagram, substitute an object or expression of the appropriate type in your SQL statement. For example, to write a `CREATE TABLE` statement, use the name of the table you want to create, such as `EMP`, in place of the *table* parameter in the syntax diagram. (Note that parameter names appear in italics in the text.)

This lists shows parameters that appear in the syntax diagrams and provides examples of the values you might substitute for them in your statements:

Parameter	Description	Examples
<i>table</i>	The substitution value must be the name of an object of the type specified by the parameter. For a list of all types of objects, see the section, " <a href="#">Schema Objects</a> " on page 2-68.	emp
<i>c</i>	The substitution value must be a single character from your database character set.	T s
<i>'text'</i>	The substitution value must be a text string in single quotes. See the syntax description of <i>'text'</i> in " <a href="#">Text</a> " on page 2-2.	'Employee records'
<i>char</i>	The substitution value must be an expression of datatype CHAR or VARCHAR2 or a character literal in single quotes.	ename 'Smith'
<i>condition</i>	The substitution value must be a condition that evaluates to TRUE or FALSE. See the syntax description of <i>condition</i> in " <a href="#">Conditions</a> " on page 5-14.	ename >'A'
<i>date</i> <i>d</i>	The substitution value must be a date constant or an expression of DATE datatype.	TO_DATE( '01-Jan-1994', 'DD-MON-YYYY')
<i>expr</i>	The substitution value can be an expression of any datatype as defined in the syntax description of <i>expr</i> in " <a href="#">Expressions</a> " on page 5-1.	sal + 1000
<i>integer</i>	The substitution value must be an integer as defined by the syntax description of integer in " <a href="#">Integer</a> " on page 2-3.	72
<i>number</i> <i>m</i> <i>n</i>	The substitution value must be an expression of NUMBER datatype or a number constant as defined in the syntax description of <i>number</i> in " <a href="#">Number</a> " on page 2-4.	AVG(sal) 15 * 7
<i>raw</i>	The substitution value must be an expression of datatype RAW.	HEXTORAW('7D')
<i>query</i>	The substitution value must be a SELECT statement that will be used in another SQL statement. See " <a href="#">SELECT and Subqueries</a> " on page 7-569.	SELECT ename FROM emp
<i>db_name</i>	The substitution value must be the name of a nondefault database in an embedded SQL program.	sales_db

Parameter	Description	Examples
<i>db_string</i>	The substitution value must be the database identification string for a Net8 database connection. For details, see the user's guide for your specific Net8 protocol.	

## Code Examples

This reference contains many examples of SQL statements. These examples show you how to use elements of SQL. The following example shows a `CREATE TABLE` statement:

```
CREATE TABLE accounts
  ( accno      NUMBER,
    owner      VARCHAR2(10),
    balance    NUMBER(7,2) );
```

Code examples appear in a different font than the text.

Examples follow these conventions:

- Keywords, such as `CREATE` and `NUMBER`, appear in uppercase.
- Names of database objects and their parts, such as `ACCOUNTS` and `ACCNO`, appear in lowercase, although they appear in uppercase in the text.
- PL/SQL blocks appear in italics. Keywords and parameters in these blocks may not be documented in this reference unless they are also SQL keywords and parameters. For more information see *PL/SQL User's Guide and Reference*.

Many examples assume the existence of objects that are not created in the example itself. The examples will not work as expected unless you first create those underlying objects.

SQL is not case sensitive (except for quoted identifiers), so you need not follow these conventions when writing your own SQL statements. However, your statements may be easier for you to read if you do.

Some Oracle tools require you to terminate SQL statements with a special character. For example, the code examples in this reference were issued through `SQL*Plus`, and therefore are terminated with a semicolon (;). If you issue these example statements to Oracle, you must terminate them with the special character expected by the Oracle tool you are using.

## Example Data

Many of the examples in this reference operate on sample tables. The definitions of some of these tables appear in a SQL script available on your distribution medium. On most operating systems the name of this script is `UTLSAMPL.SQL`, although its exact name and location depend on your operating system. This script creates sample users and creates these sample tables in the schema of the user `SCOTT` (password `TIGER`):

```
CREATE TABLE dept
  (deptno    NUMBER(2)          CONSTRAINT pk_dept PRIMARY KEY,
   dname     VARCHAR2(14),
   loc       VARCHAR2(13) );

CREATE TABLE emp
  (empno     NUMBER(4)          CONSTRAINT pk_emp PRIMARY KEY,
   ename     VARCHAR2(10),
   job       VARCHAR2(9),
   mgr       NUMBER(4),
   hiredate  DATE,
   sal       NUMBER(7,2),
   comm      NUMBER(7,2),
   deptno    NUMBER(2)          CONSTRAINT fk_deptno REFERENCES dept );

CREATE TABLE bonus
  (ename     VARCHAR2(10),
   job       VARCHAR2(9),
   sal       NUMBER,
   comm      NUMBER );

CREATE TABLE salgrade
  (grade     NUMBER,
   losal     NUMBER,
   hisal     NUMBER );
```

The script also fills the sample tables with this data:

```
SELECT * FROM dept;
```

DEPTNO	DNAME	LOC
10	ACCOUNTING	NEW YORK
20	RESEARCH	DALLAS
30	SALES	CHICAGO
40	OPERATIONS	BOSTON

```
SELECT * FROM emp;
```

EMPNO	ENAME	JOB	MGR	HIREDATE	SAL	COMM	DEPTNO
-------	-------	-----	-----	----------	-----	------	--------



```

-----
7369 SMITH CLERK 7902 17-DEC-80 800 20
7499 ALLEN SALESMAN 7698 20-FEB-81 1600 300 30
7521 WARD SALESMAN 7698 22-FEB-81 1250 500 30
7566 JONES MANAGER 7839 02-APR-81 2975 20
7654 MARTIN SALESMAN 7698 28-SEP-81 1250 1400 30
7698 BLAKE MANAGER 7839 01-MAY-81 2850 30
7782 CLARK MANAGER 7839 09-JUN-81 2450 10
7788 SCOTT ANALYST 7566 19-APR-87 3000 20
7839 KING PRESIDENT 17-NOV-81 5000 10
7844 TURNER SALESMAN 7698 08-SEP-81 1500 30
7876 ADAMS CLERK 7788 23-MAY-87 1100 20
7900 JAMES CLERK 7698 03-DEC-81 950 30
7902 FORD ANALYST 7566 03-DEC-81 3000 20
7934 MILLER CLERK 7782 23-JAN-82 1300 10

```

```
SELECT * FROM salgrade;
```

```

GRADE LOSAL HISAL
-----
1      700    1200
2      1201   1400
3      1401   2000
4      2001   3000
5      3001   9999

```

To perform all the operations of the script, run it when you are logged into Oracle as the user SYSTEM.

## Your Comments Are Welcome

We value and appreciate your comments as an Oracle user and reader of our references. As we write, revise, and evaluate, your opinions are the most important input we receive. At the front of this reference is a reader's comment form that we encourage you to use to tell us both what you like and what you dislike about this (or other) Oracle manuals. If the form is missing, or you would like to contact us, please use the following address or fax number:

Server Technologies Documentation Manager  
Oracle Corporation  
500 Oracle Parkway  
Redwood City, CA 94065  
FAX: 650-506-7228

You can also e-mail your comments to the Information Development department at the following e-mail address: [infodev@us.oracle.com](mailto:infodev@us.oracle.com)

---

# Introduction

*The chief merit of language is clearness . . . .*

*Galen, On the Natural Faculties*

Structured Query Language (SQL) is the set of statements with which all programs and users access data in an Oracle database. Application programs and Oracle tools often allow users access to the database without using SQL directly, but these applications in turn must use SQL when executing the user's request. This chapter provides background information on SQL as used by most relational database systems. Topics include:

- [History of SQL](#)
- [SQL Standards](#)
- [Embedded SQL](#)
- [Lexical Conventions](#)
- [Tools Support](#)

## History of SQL

Dr. E. F. Codd published the paper, "A Relational Model of Data for Large Shared Data Banks", in June 1970 in the Association of Computer Machinery (ACM) journal, *Communications of the ACM*. Codd's model is now accepted as the definitive model for relational database management systems (RDBMS). The language, Structured English Query Language ("SEQUEL") was developed by IBM Corporation, Inc., to use Codd's model. SEQUEL later became SQL (still pronounced "sequel"). In 1979, Relational Software, Inc. (now Oracle Corporation) introduced the first commercially available implementation of SQL. Today, SQL is accepted as the standard RDBMS language.

---

## SQL Standards

Oracle Corporation strives to comply with industry-accepted standards and participates actively in SQL standards committees. Industry-accepted committees are the American National Standards Institute (ANSI) and the International Standards Organization (ISO), which is affiliated with the International Electrotechnical Commission (IEC). Both ANSI and the ISO/IEC have accepted SQL as the standard language for relational databases. When a new SQL standard is simultaneously published by these organizations, the names of the standards conform to conventions used by the organization, but the standards are technically identical.

The latest SQL standard was adopted in July 1999 and is often called SQL-99. The formal names of this standard are:

- ANSI X3.135-1999, "Database Language SQL", Parts 1 ("Framework"), 2 ("Foundation"), and 5 ("Bindings")
- ISO/IEC 9075:1999, "Database Language SQL", Parts 1 ("Framework"), 2 ("Foundation"), and 5 ("Bindings")

SQL-99 replaced the previous version of the standard, commonly known as SQL-92. SQL-99 is an upward compatible extension of SQL-92, except for a few minor incompatibilities noted in Annex E of Part 2, "Foundation," of SQL-99.

SQL-92 defined four levels of compliance: Entry, Transitional, Intermediate, and Full. A conforming SQL implementation must support at least Entry SQL. Oracle8i fully supports Entry SQL as outlined in Federal Information Processing Standard (FIPS) PUB 127-2, and has many features that conform to Transitional, Intermediate, or Full SQL.

The minimal conformance level for SQL-99 is known as Core. Core SQL-99 is a superset of SQL-92 Entry Level specification. Oracle8i also is broadly compatible with the SQL-99 Core specification. However, some SQL-99 Core features are not currently implemented in Oracle8i or differ from the Oracle8i implementation. Oracle Corporation is committed to fully supporting SQL-99 Core functionality in a future release, while providing upward compatibility for existing applications.

**See Also:** [Appendix B, "Oracle and Standard SQL"](#) for more information about Oracle and standard SQL.

### How SQL Works

The strengths of SQL provide benefits for all types of users, including application programmers, database administrators, managers, and end users. Technically

---

speaking, SQL is a data sublanguage. The purpose of SQL is to provide an interface to a relational database such as Oracle, and all SQL statements are instructions to the database. In this SQL differs from general-purpose programming languages like C and BASIC. Among the features of SQL are the following:

- It processes sets of data as groups rather than as individual units.
- It provides automatic navigation to the data.
- It uses statements that are complex and powerful individually, and that therefore stand alone. Flow-control statements were not part of SQL originally, but they are found in the recently accepted optional part of SQL, ISO/IEC 9075-5: 1996. Flow-control statements are commonly known as "persistent stored modules" (PSM), and Oracle's PL/SQL extension to SQL is similar to PSM.

Essentially, SQL lets you work with data at the logical level. You need to be concerned with the implementation details only when you want to manipulate the data. For example, to retrieve a set of rows from a table, you define a condition used to filter the rows. All rows satisfying the condition are retrieved in a single step and can be passed as a unit to the user, to another SQL statement, or to an application. You need not deal with the rows one by one, nor do you have to worry about how they are physically stored or retrieved. All SQL statements use the **optimizer**, a part of Oracle that determines the most efficient means of accessing the specified data. Oracle also provides techniques that you can use to make the optimizer perform its job better.

SQL provides statements for a variety of tasks, including:

- Querying data
- Inserting, updating, and deleting rows in a table
- Creating, replacing, altering, and dropping objects
- Controlling access to the database and its objects
- Guaranteeing database consistency and integrity

SQL unifies all of the above tasks in one consistent language.

### **Common Language for All Relational Databases**

All major relational database management systems support SQL, so you can transfer all skills you have gained with SQL from one database to another. In addition, all programs written in SQL are portable. They can often be moved from one database to another with very little modification.

## Embedded SQL

Embedded SQL refers to the use of standard SQL statements embedded within a procedural programming language. The embedded SQL statements are documented in the Oracle precompiler books: *Pro\*C/C++ Precompiler Programmer's Guide*, *Pro\*COBOL Precompiler Programmer's Guide*, and *SQL\*Module for Ada Programmer's Guide*.

Embedded SQL is a collection of these statements:

- All SQL commands, such as `SELECT` and `INSERT`, available with SQL with interactive tools
- Dynamic SQL execution commands, such as `PREPARE` and `OPEN`, which integrate the standard SQL statements with a procedural programming language

Embedded SQL also includes extensions to some standard SQL statements. Embedded SQL is supported by the Oracle precompilers. The Oracle precompilers interpret embedded SQL statements and translate them into statements that can be understood by procedural language compilers.

Each of these Oracle precompilers translates embedded SQL programs into a different procedural language:

- Pro\*C/C++ precompiler
- Pro\*COBOL precompiler
- SQL\*Module for ADA

**See Also:** For a definition of the Oracle precompilers and the embedded SQL statements, see *SQL\*Module for Ada Programmer's Guide*, *Pro\*C/C++ Precompiler Programmer's Guide*, and *Pro\*COBOL Precompiler Programmer's Guide*.

## Lexical Conventions

The following lexical conventions for issuing SQL statements apply specifically to Oracle's implementation of SQL, but are generally acceptable in other SQL implementations.

When you issue a SQL statement, you can include one or more tabs, carriage returns, spaces, or comments anywhere a space occurs within the definition of the statement. Thus, Oracle evaluates the following two statements in the same manner:

```
SELECT  ENAME , SAL*12 , MONTHS_BETWEEN( HIREDATE , SYSDATE) FROM EMP ;
```

```
SELECT ENAME,  
       SAL * 12,  
       MONTHS_BETWEEN( HIREDATE, SYSDATE )  
FROM EMP;
```

Case is insignificant in reserved words, keywords, identifiers and parameters. However, case is significant in text literals and quoted names. See the syntax description in ["Text"](#) on page 2-2.

## Tools Support

Most (but not all) Oracle tools support all features of Oracle SQL. This reference describes the complete functionality of SQL. If the Oracle tool that you are using does not support this complete functionality, you can find a discussion of the restrictions in the manual describing the tool, such as *SQL\*Plus User's Guide and Reference*.

If you are using Trusted Oracle, see your Trusted Oracle documentation for information about SQL statements specific to that environment.





---

---

# Basic Elements of Oracle SQL

*Once the whole is divided, the parts need names.*

Lao Tsu, *Tao Te Ching: Thirty-Two*

This chapter contains reference information on the basic elements of Oracle SQL. These elements are simplest building blocks of SQL statements. Therefore, before using the statements described in [Chapter 7, "SQL Statements"](#), you should familiarize yourself with the concepts covered in this chapter, as well as in [Chapter 3, "Operators"](#), [Chapter 4, "Functions"](#), [Chapter 5, "Expressions, Conditions, and Queries"](#), and [Chapter 6, "About SQL Statements"](#):

- Literals
  - Text
  - Integer
  - Number
  - Interval
- Datatypes
- Format Models
- Nulls
- Pseudocolumns
- Comments
- Database Objects
- Schema Object Names and Qualifiers
- Referring to Schema Objects and Parts

## Literals

The terms **literal** and **constant value** are synonymous and refer to a fixed data value. For example, 'JACK', 'BLUE ISLAND', and '101' are all character literals; 5001 is a numeric literal. Note that character literals are enclosed in single quotation marks, which enable Oracle to distinguish them from schema object names.


Many SQL statements and functions require you to specify character and numeric literal values. You can also specify literals as part of expressions and conditions. You can specify character literals with the *'text'* notation, national character literals with the *N'text'* notation, and numeric literals with the *integer* or *number* notation, depending on the context of the literal. The syntactic forms of these notations appear in the sections that follow.

## Text

Text specifies a text or character literal. You must use this notation to specify values whenever *'text'* or *char* appear in expressions, conditions, SQL functions, and SQL statements in other parts of this reference.

The syntax of text is as follows:

**text::=**



where

- N** specifies representation of the literal using the national character set. Text entered using this notation is translated into the national character set by Oracle when used.
- c** is any member of the user's character set, except a single quotation mark (').
- ' '** are two single quotation marks that begin and end text literals. To represent one single quotation mark within a literal, enter two single quotation marks.

A text literal must be enclosed in single quotation marks. This reference uses the terms **text literal** and **character literal** interchangeably.

Text literals have properties of both the `CHAR` and `VARCHAR2` datatypes:

- Within expressions and conditions, Oracle treats text literals as though they have the datatype `CHAR` by comparing them using blank-padded comparison semantics. See "[Blank-Padded Comparison Semantics](#)" on page 2-33.
- A text literal can have a maximum length of 4000 bytes.

Here are some valid text literals:

```
'Hello'
'ORACLE.dbs'
'Jackie''s raincoat'
'09-MAR-98'
N'nchar literal'
```


**See Also:** "[Expressions](#)" on page 5-1 for the syntax description of *expr*.

## Integer

You must use the integer notation to specify an integer whenever *integer* appears in expressions, conditions, SQL functions, and SQL statements described in other parts of this reference.

The syntax of *integer* is as follows:

**integer::=**



where

*digit* is one of 0, 1, 2, 3, 4, 5, 6, 7, 8, 9.

An integer can store a maximum of 38 digits of precision.

Here are some valid integers:

```
7
+255
```

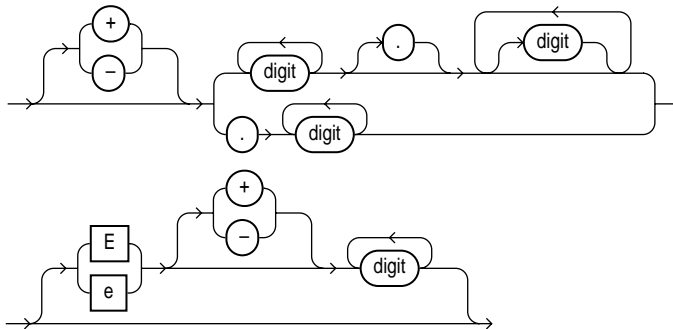
**See Also:** ["Expressions"](#) on page 5-1 for the syntax description of *expr*.

## Number

You must use the number notation to specify values whenever *number* appears in expressions, conditions, SQL functions, and SQL statements in other parts of this reference.

The syntax of *number* is as follows:

**number::=**



where

- +, -** indicates a positive or negative value. If you omit the sign, a positive value is the default.
- digit** is one of 0, 1, 2, 3, 4, 5, 6, 7, 8 or 9.
- e, E** indicates that the number is specified in scientific notation. The digits after the E specify the exponent. The exponent can range from -130 to 125.

A *number* can store a maximum of 38 digits of precision.

If you have established a decimal character other than a period (.) with the initialization parameter `NLS_NUMERIC_CHARACTERS`, you must specify numeric literals with *'text'* notation. In such cases, Oracle automatically converts the text literal to a numeric value.

For example, if the `NLS_NUMERIC_CHARACTERS` parameter specifies a decimal character of comma, specify the number 5.123 as follows:

```
'5,123'
```

**See Also:** ["ALTER SESSION"](#) on page 7-83 and *Oracle8i Reference*.

Here are some valid representations of *number*:

25  
+6.34  
0.5  
25e-03  
-1

**See Also:** ["Expressions"](#) on page 5-1 for the syntax description of *expr*.

## Interval

An interval literal specifies a period of time. You can specify these differences in terms of years and months, or in terms of days, hours, minutes, and seconds. Oracle supports two types of interval literals, YEAR TO MONTH and DAY TO SECOND. Each type contains a leading field and may contain a trailing field. The leading field defines the basic unit of date or time being measured. The trailing field defines the smallest increment of the basic unit being considered. For example, a YEAR TO MONTH interval considers an interval of years to the nearest month. A DAY TO MINUTE interval considers an interval of days to the nearest minute.

If you have date data in numeric form, you can use the NUMTOYMINTERVAL or NUMTODSINTERVAL conversion function to convert the numeric data into interval literals.


Interval literals are used primarily with analytic functions.

**See Also:**

- ["Analytic Functions"](#) on page 4-7 and *Oracle8i Data Warehousing Guide*.
- ["NUMTODSINTERVAL"](#) on page 4-62 and ["NUMTOYMINTERVAL"](#) on page 4-63.

## INTERVAL YEAR TO MONTH

Specify YEAR TO MONTH interval literals using the following syntax:



where

- *'integer [-integer]'* specifies integer values for the leading and optional trailing field of the literal. If the leading field is YEAR and the trailing field is MONTH, the range of integer values for the month field is 0 to 11.
- *precision* is the number of digits in the leading field. The valid range of the leading field precision is 0 to 9 and its default value is 2.

**Restriction:** The leading field must be a larger time element than the trailing field. For example, INTERVAL '0-1' MONTH TO YEAR is not valid.

The following INTERVAL YEAR TO MONTH literal indicates an interval of 123 years, 2 months:

```
INTERVAL '123-2' YEAR(3) TO MONTH
```

Examples of the other forms of the literal follow, including some abbreviated versions:

```
INTERVAL '123-2' YEAR(3) TO MONTH
```

indicates an interval of 123 years, 2 months. You must specify the leading field precision if it is greater than the default of 2 digits.

```
INTERVAL '123' YEAR(3)
```

indicates an interval of 123 years 0 months.

```
INTERVAL '300' MONTH(3)
```

indicates an interval of 300 months.

```
INTERVAL '4' YEAR
```

maps to INTERVAL '4-0' YEAR TO MONTH and indicates 4 years.

INTERVAL '50' MONTH

maps to INTERVAL '4-2' YEAR TO MONTH and indicates 50 months or 4 years 2 months.

INTERVAL '123' YEAR

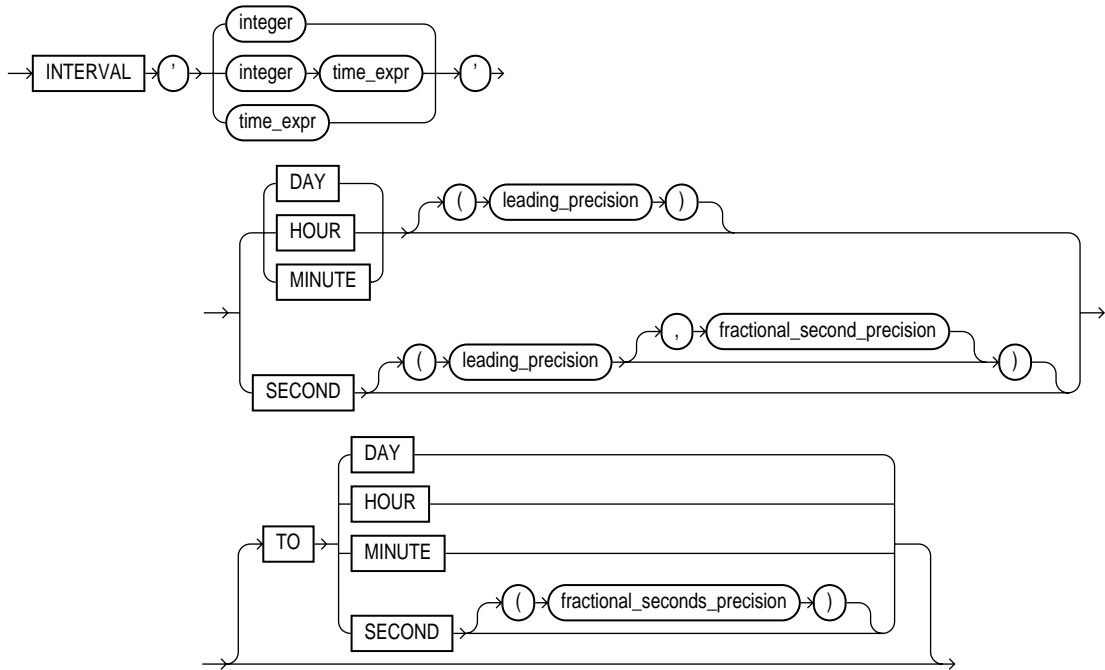
returns an error, because the default precision is 2, and '123' has 3 digits.

You can add or subtract one INTERVAL YEAR TO MONTH literal to or from another to yield another INTERVAL YEAR TO MONTH literal. For example:

INTERVAL '5-3' YEAR TO MONTH + INTERVAL '20' MONTH TO MONTH =  
INTERVAL '6-11' YEAR TO MONTH

### INTERVAL DAY TO SECOND

Specify DAY TO SECOND interval literals using the following syntax:



where

- *integer* specifies the number of days. If this value contains more digits than the number specified by the leading precision, Oracle returns an error.

- *time\_expr* specifies a time in the format HH[ :MI[ :SS[ .n] ] ] or MI[ :SS[ .n] ] or SS[ .n], where *n* specifies the fractional part of a second. If *n* contains more digits than the number specified by *fractional\_seconds\_precision*, then *n* is rounded to the number of digits specified by the *fractional\_seconds\_precision* value. You can specify *time\_expr* following an integer and a space only if the leading field is DAY.
- *day\_precision* is the number of digits in the DAY date field. Accepted values are 0 to 9. The default is 2.
- *fractional\_seconds\_precision* is the number of digits in the fractional part of the SECOND datetime field. Accepted values are 1 to 9. The default is 6.

**Restriction:** The leading field must be a larger time element than the trailing field. For example, INTERVAL MINUTE TO DAY is not valid. As a result of this restriction, if SECOND is the leading field, the interval literal cannot have any trailing field.

The valid range of values for the trailing field are as follows:

HOUR	0 to 23
MINUTE	0 to 59
SECOND	0 to 59.999999999

Examples of the various forms of INTERVAL DAY TO SECOND literals follow, including some abbreviated versions:

INTERVAL '4 5:12:10.222' DAY(3) TO SECOND(3)	indicates 4 days, 5 hours, 12 minutes, 10 seconds, and 222 thousandths of a second.
INTERVAL '4 5:12' DAY TO MINUTE	indicates 4 days, 5 hours and 12 minutes.
INTERVAL '400 5' DAY(3) TO HOUR	indicates 400 days 5 hours.
INTERVAL '400' DAY(3)	indicates 400 days.
INTERVAL '11:12:10.2222222' HOUR TO SECOND(7)	indicates 11 hours, 12 minutes, and 10.2222222 seconds.
INTERVAL '11:20' HOUR TO MINUTE	indicates 11 hours and 20 minutes.
INTERVAL '10' HOUR	indicates 10 hours.
INTERVAL '10:22' MINUTE TO SECOND	indicates 10 minutes 22 seconds.



<code>INTERVAL '10' MINUTE</code>	indicates 10 minutes.
<code>INTERVAL '4' DAY</code>	indicates 4 days.
<code>INTERVAL '25' HOUR</code>	indicates 25 hours.
<code>INTERVAL '40' MINUTE</code>	indicates 40 minutes.
<code>INTERVAL '120' HOUR(3)</code>	indicates 120 hours
<code>INTERVAL '30.12345' SECOND(2,4)</code>	indicates 30.1235 seconds. The fractional second '12345' is rounded to '1235' because the precision is 4.

You can add or subtract one `DAY TO SECOND` interval literal from another `DAY TO SECOND` literal. For example,

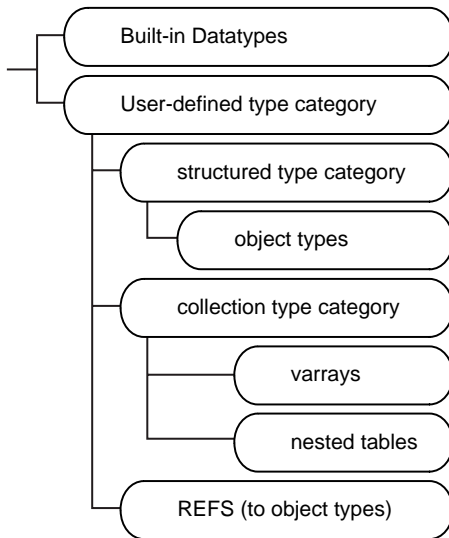
```
INTERVAL '20' DAY - INTERVAL '240' HOUR = INTERVAL '10' DAY
```

## Datatypes

Each literal or column value manipulated by Oracle has a **datatype**. A value's datatype associates a fixed set of properties with the value. These properties cause Oracle to treat values of one datatype differently from values of another. For example, you can add values of `NUMBER` datatype, but not values of `RAW` datatype.

When you create a table or cluster, you must specify a datatype for each of its columns. When you create a procedure or stored function, you must specify a datatype for each of its arguments. These datatypes define the domain of values that each column can contain or each argument can have. For example, `DATE` columns cannot accept the value February 29 (except for a leap year) or the values 2 or 'SHOE'. Each value subsequently placed in a column assumes the column's datatype. For example, if you insert '01-JAN-98' into a `DATE` column, Oracle treats the '01-JAN-98' character string as a `DATE` value after verifying that it translates to a valid date.

Oracle provides a number of built-in datatypes as well as several categories for user-defined types, as shown in [Figure 2-1](#).

**Figure 2–1 Oracle Type Categories**

The syntax of the Oracle built-in datatypes appears in the next diagram. [Table 2–1](#) summarizes Oracle built-in datatypes. The rest of this section describes these datatypes as well as the various kinds of user-defined types.

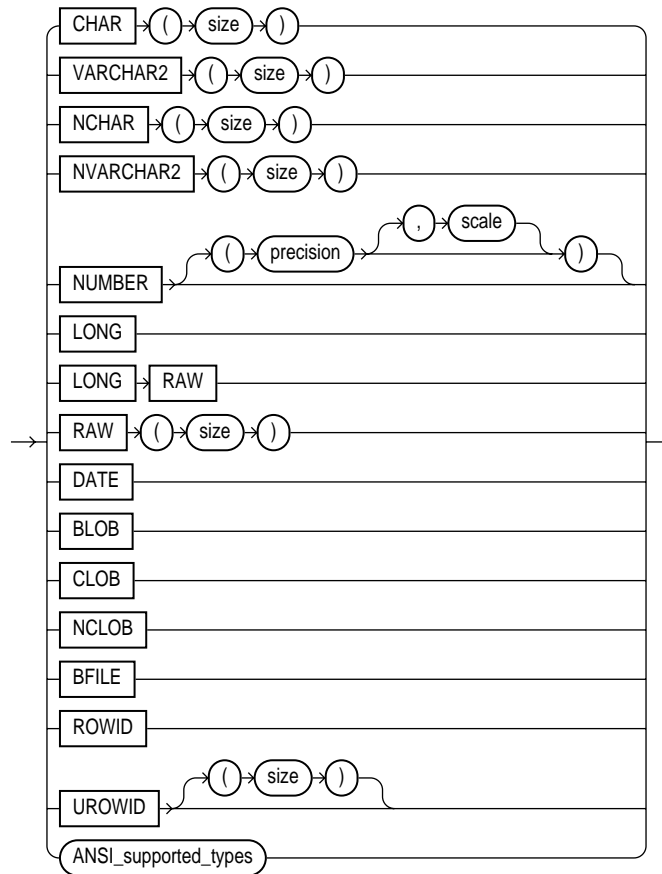
---

---

**Note:** The Oracle precompilers recognize other datatypes in embedded SQL programs. These datatypes are called *external datatypes* and are associated with host variables. Do not confuse built-in and user-defined datatypes with external datatypes. For information on external datatypes, including how Oracle converts between them and built-in or user-defined datatypes, see *Pro\*COBOL Precompiler Programmer's Guide*, *Pro\*C/C++ Precompiler Programmer's Guide*, and *SQL\*Module for Ada Programmer's Guide*.

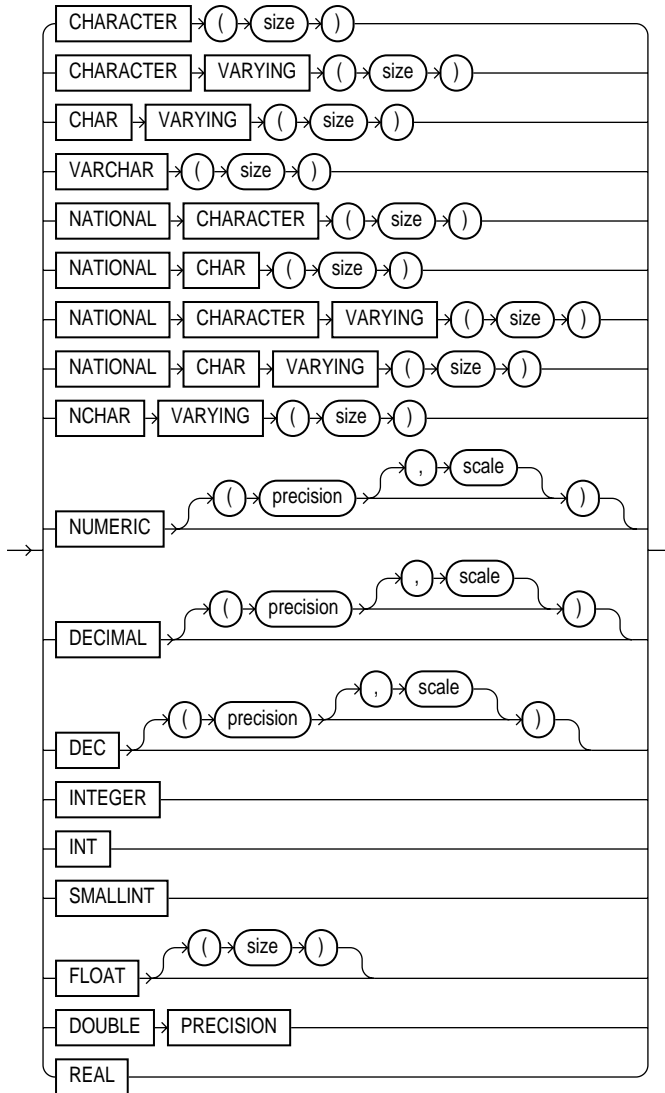
---

---

**built-in datatypes:**

The ANSI-supported datatypes appear in the figure that follows. [Table 2-2](#) shows the mapping of ANSI-supported datatypes to Oracle build-in datatypes.

**ANSI-supported datatypes:**



**Table 2–1 Built-In Datatype Summary**

Code <sup>a</sup>	Built-In Datatype	Description
1	VARCHAR2( <i>size</i> )	Variable-length character string having maximum length <i>size</i> bytes. Maximum <i>size</i> is 4000, and minimum is 1. You must specify <i>size</i> for VARCHAR2.
1	NVARCHAR2( <i>size</i> )	Variable-length character string having maximum length <i>size</i> characters or bytes, depending on the choice of national character set. Maximum <i>size</i> is determined by the number of bytes required to store each character, with an upper limit of 4000 bytes. You must specify <i>size</i> for NVARCHAR2.
2	NUMBER( <i>p</i> , <i>s</i> )	Number having precision <i>p</i> and scale <i>s</i> . The precision <i>p</i> can range from 1 to 38. The scale <i>s</i> can range from -84 to 127.
8	LONG	Character data of variable length up to 2 gigabytes, or $2^{31} - 1$ bytes.
12	DATE	Valid date range from January 1, 4712 BC to December 31, 9999 AD.
23	RAW( <i>size</i> )	Raw binary data of length <i>size</i> bytes. Maximum <i>size</i> is 2000 bytes. You must specify <i>size</i> for a RAW value.
24	LONG RAW	Raw binary data of variable length up to 2 gigabytes.
69	ROWID	Hexadecimal string representing the unique address of a row in its table. This datatype is primarily for values returned by the ROWID pseudocolumn.
208	UROWID [( <i>size</i> )]	Hexadecimal string representing the logical address of a row of an index-organized table. The optional <i>size</i> is the size of a column of type UROWID. The maximum size and default is 4000 bytes.
96	CHAR( <i>size</i> )	Fixed-length character data of length <i>size</i> bytes. Maximum <i>size</i> is 2000 bytes. Default and minimum <i>size</i> is 1 byte.

<sup>a</sup> The codes listed for the datatypes are used internally by Oracle. The datatype code of a column or object attribute is returned when you use the DUMP function.

**Table 2–1 (Cont.) Built-In Datatype Summary**

Code <sup>a</sup>	Built-In Datatype	Description
96	NCHAR( <i>size</i> )	Fixed-length character data of length <i>size</i> characters or bytes, depending on the choice of national character set. Maximum <i>size</i> is determined by the number of bytes required to store each character, with an upper limit of 2000 bytes. Default and minimum <i>size</i> is 1 character or 1 byte, depending on the character set.
112	CLOB	A character large object containing single-byte characters. Both fixed-width and variable-width character sets are supported, both using the CHAR database character set. Maximum size is 4 gigabytes.
112	NCLOB	A character large object containing multibyte characters. Both fixed-width and variable-width character sets are supported, both using the NCHAR database character set. Maximum size is 4 gigabytes. Stores national character set data.
113	BLOB	A binary large object. Maximum size is 4 gigabytes.
114	BFILE	Contains a locator to a large binary file stored outside the database. Enables byte stream I/O access to external LOBs residing on the database server. Maximum size is 4 gigabytes.

<sup>a</sup> The codes listed for the datatypes are used internally by Oracle. The datatype code of a column or object attribute is returned when you use the DUMP function.

## Character Datatypes

Character datatypes store character (alphanumeric) data, which are words and free-form text, in the database character set or national character set. They are less restrictive than other datatypes and consequently have fewer properties. For example, character columns can store all alphanumeric values, but NUMBER columns can store only numeric values.

Character data is stored in strings with byte values corresponding to one of the character sets, such as 7-bit ASCII or EBCDIC Code Page 500, specified when the database was created. Oracle supports both single-byte and multibyte character sets.

These datatypes are used for character data:

- [CHAR Datatype](#)

- [NCHAR Datatype](#)
- [NVARCHAR2 Datatype](#)
- [VARCHAR2 Datatype](#)

## CHAR Datatype

The `CHAR` datatype specifies a fixed-length character string. When you create a table with a `CHAR` column, you supply the column length in bytes. Oracle subsequently ensures that all values stored in that column have this length. If you insert a value that is shorter than the column length, Oracle blank-pads the value to column length. If you try to insert a value that is too long for the column, Oracle returns an error.

The default length for a `CHAR` column is 1 character and the maximum allowed is 2000 characters. A zero-length string can be inserted into a `CHAR` column, but the column is blank-padded to 1 character when used in comparisons. For information on comparison semantics, see "[Datatype Comparison Rules](#)" on page 2-32.

---

---

**Note:** To ensure proper data conversion between databases with different character sets, you must ensure that `CHAR` data consists of well-formed strings. See *Oracle8i National Language Support Guide* for more information on character set support.

---

---

## NCHAR Datatype

The `NCHAR` datatype specifies a fixed-length national character set character string. When you create a table with an `NCHAR` column, you define the column length either in characters or in bytes. You define the national character set when you create your database.

If the national character set of the database is fixed width, such as `JA16EUCFIXED`, then you declare the `NCHAR` column size as the number of characters desired for the string length. If the national character set is variable width, such as `JA16SJIS`, you declare the column size in bytes. The following statement creates a table with one `NCHAR` column that can store strings up to 30 characters in length using `JA16EUCFIXED` as the national character set:

```
CREATE TABLE tabl (coll NCHAR(30));
```

The column's maximum length is determined by the national character set definition. Width specifications of character datatype `NCHAR` refer to the number of characters if the national character set is fixed width and refer to the number of

bytes if the national character set is variable width. The maximum column size allowed is 2000 bytes. For fixed-width, multibyte character sets, the maximum length of a column allowed is the number of characters that fit into no more than 2000 bytes.

If you insert a value that is shorter than the column length, Oracle blank-pads the value to column length. You cannot insert a CHAR value into an NCHAR column, nor can you insert an NCHAR value into a CHAR column.

The following example compares the COL1 column of TAB1 with national character set string 'NCHAR literal':

```
SELECT * FROM tab1 WHERE col1 = N'NCHAR literal';
```

### **NVARCHAR2 Datatype**

The NVARCHAR2 datatype specifies a variable-length national character set character string. When you create a table with an NVARCHAR2 column, you supply the maximum number of characters or bytes it can hold. Oracle subsequently stores each value in the column exactly as you specify it, provided the value does not exceed the column's maximum length.

The column's maximum length is determined by the national character set definition. Width specifications of character datatype NVARCHAR2 refer to the number of characters if the national character set is fixed width and refer to the number of bytes if the national character set is variable width. The maximum column size allowed is 4000 bytes. For fixed-width, multibyte character sets, the maximum length of a column allowed is the number of characters that fit into no more than 4000 bytes.

The following statement creates a table with one NVARCHAR2 column of 2000 characters in length (stored as 4000 bytes, because each character takes two bytes) using JA16EUCFIXED as the national character set:

```
CREATE TABLE tab1 (col1 NVARCHAR2(2000));
```

### **VARCHAR2 Datatype**

The VARCHAR2 datatype specifies a variable-length character string. When you create a VARCHAR2 column, you supply the maximum number of bytes of data that it can hold. Oracle subsequently stores each value in the column exactly as you specify it, provided the value does not exceed the column's maximum length. If you try to insert a value that exceeds the specified length, Oracle returns an error.

You must specify a maximum length for a VARCHAR2 column. This maximum must be at least 1 byte, although the actual length of the string stored is permitted to be



zero. The maximum length of VARCHAR2 data is 4000 bytes. Oracle compares VARCHAR2 values using nonpadded comparison semantics. For information on comparison semantics, see "[Datatype Comparison Rules](#)" on page 2-32.

---

---

**Note:** To ensure proper data conversion between databases with different character sets, you must ensure that VARCHAR2 data consists of well-formed strings. See *Oracle8i National Language Support Guide* for more information on character set support.

---

---

## VARCHAR Datatype

The VARCHAR datatype is currently synonymous with the VARCHAR2 datatype. Oracle recommends that you use VARCHAR2 rather than VARCHAR. In the future, VARCHAR might be defined as a separate datatype used for variable-length character strings compared with different comparison semantics.

## NUMBER Datatype

The NUMBER datatype stores zero, positive, and negative fixed and floating-point numbers with magnitudes between  $1.0 \times 10^{-130}$  and  $9.9...9 \times 10^{125}$  (38 nines followed by 88 zeroes) with 38 digits of precision. If you specify an arithmetic expression whose value has a magnitude greater than or equal to  $1.0 \times 10^{126}$ , Oracle returns an error.

Specify a fixed-point number using the following form:

`NUMBER(p, s)`

where:

- p* is the *precision*, or the total number of digits. Oracle guarantees the portability of numbers with precision ranging from 1 to 38.
- s* is the *scale*, or the number of digits to the right of the decimal point. The scale can range from -84 to 127.

Specify an integer using the following form:

`NUMBER(p)` is a fixed-point number with precision *p* and scale 0. This is equivalent to `NUMBER(p, 0)`.

Specify a floating-point number using the following form:

`NUMBER` is a floating-point number with decimal precision 38. Note that a scale value is not applicable for floating-point numbers. (See ["Floating-Point Numbers"](#) on page 2-19 for more information.)

### Scale and Precision

Specify the scale and precision of a fixed-point number column for extra integrity checking on input. Specifying scale and precision does not force all values to a fixed length. If a value exceeds the precision, Oracle returns an error. If a value exceeds the scale, Oracle rounds it.

The following examples show how Oracle stores data using different precisions and scales.

<b>Actual Data</b>	<b>Specified As</b>	<b>Stored As</b>
7456123.89	<code>NUMBER</code>	7456123.89
7456123.89	<code>NUMBER ( 9 )</code>	7456124
7456123.89	<code>NUMBER ( 9 , 2 )</code>	7456123.89
7456123.89	<code>NUMBER ( 9 , 1 )</code>	7456123.9
7456123.89	<code>NUMBER ( 6 )</code>	exceeds precision
7456123.89	<code>NUMBER ( 7 , - 2 )</code>	7456100
7456123.89	<code>NUMBER ( - 7 , 2 )</code>	exceeds precision

### Negative Scale

If the scale is negative, the actual data is rounded to the specified number of places to the left of the decimal point. For example, a specification of (10,-2) means to round to hundreds.

### Scale Greater than Precision

You can specify a scale that is greater than precision, although it is uncommon. In this case, the precision specifies the maximum number of digits to the right of the decimal point. As with all number datatypes, if the value exceeds the precision, Oracle returns an error message. If the value exceeds the scale, Oracle rounds the value. For example, a column defined as `NUMBER(4,5)` requires a zero for the first digit after the decimal point and rounds all values past the fifth digit after the decimal point. The following examples show the effects of a scale greater than precision:

Actual Data	Specified As	Stored As
.01234	NUMBER(4,5)	.01234
.00012	NUMBER(4,5)	.00012
.000127	NUMBER(4,5)	.00013
.0000012	NUMBER(2,7)	.0000012
.00000123	NUMBER(2,7)	.0000012

### Floating-Point Numbers

Oracle allows you to specify floating-point numbers, which can have a decimal point anywhere from the first to the last digit or can have no decimal point at all. A scale value is not applicable to floating-point numbers, because the number of digits that can appear after the decimal point is not restricted.

You can specify floating-point numbers with the form discussed in "[NUMBER Datatype](#)" on page 2-17. Oracle also supports the ANSI datatype `FLOAT`. You can specify this datatype using one of these syntactic forms:

<code>FLOAT</code>	specifies a floating-point number with decimal precision 38, or binary precision 126.
<code>FLOAT(<i>b</i>)</code>	specifies a floating-point number with binary precision <i>b</i> . The precision <i>b</i> can range from 1 to 126. To convert from binary to decimal precision, multiply <i>b</i> by 0.30103. To convert from decimal to binary precision, multiply the decimal precision by 3.32193. The maximum of 126 digits of binary precision is roughly equivalent to 38 digits of decimal precision.

### LONG Datatype

`LONG` columns store variable-length character strings containing up to 2 gigabytes, or  $2^{31}-1$  bytes. `LONG` columns have many of the characteristics of `VARCHAR2` columns. You can use `LONG` columns to store long text strings. The length of `LONG` values may be limited by the memory available on your computer.

---

**Note:** Oracle Corporation strongly recommends that you convert `LONG` columns to `LOB` columns. `LOB` columns are subject to far fewer restrictions than `LONG` columns. See "[TO\\_LOB](#)" on page 4-102 for more information.

---

You can reference `LONG` columns in SQL statements in these places:

- `SELECT` lists
- `SET` clauses of `UPDATE` statements
- `VALUES` clauses of `INSERT` statements

The use of `LONG` values is subject to some restrictions:

- A table cannot contain more than one `LONG` column.
- You cannot create an object type with a `LONG` attribute.
- `LONG` columns cannot appear in integrity constraints (except for `NULL` and `NOT NULL` constraints).
- `LONG` columns cannot be indexed.
- A stored function cannot return a `LONG` value.
- Within a single SQL statement, all `LONG` columns, updated tables, and locked tables must be located on the same database.

`LONG` columns cannot appear in certain parts of SQL statements:

- `WHERE` clauses, `GROUP BY` clauses, `ORDER BY` clauses, or `CONNECT BY` clauses or with the `DISTINCT` operator in `SELECT` statements
- The `UNIQUE` operator of a `SELECT` statement
- The column list of a `CREATE CLUSTER` statement
- The `CLUSTER` clause of a `CREATE MATERIALIZED VIEW` statement
- SQL functions (such as `SUBSTR` or `INSTR`)
- Expressions or conditions
- `SELECT` lists of queries containing `GROUP BY` clauses
- `SELECT` lists of subqueries or queries combined by the `UNION`, `INTERSECT`, or `MINUS` set operators
- `SELECT` lists of `CREATE TABLE ... AS SELECT` statements
- `SELECT` lists in subqueries in `INSERT` statements

Triggers can use the `LONG` datatype in the following manner:

- A SQL statement within a trigger can insert data into a `LONG` column.

- If data from a `LONG` column can be converted to a constrained datatype (such as `CHAR` and `VARCHAR2`), a `LONG` column can be referenced in a SQL statement within a trigger.
- Variables in triggers cannot be declared using the `LONG` datatype.
- `:NEW` and `:OLD` cannot be used with `LONG` columns.

You can use the Oracle Call Interface functions to retrieve a portion of a `LONG` value from the database. See *Oracle Call Interface Programmer's Guide*.

## DATE Datatype

The `DATE` datatype stores date and time information. Although date and time information can be represented in both `CHAR` and `NUMBER` datatypes, the `DATE` datatype has special associated properties. For each `DATE` value, Oracle stores the following information: century, year, month, day, hour, minute, and second.

To specify a date value, you must convert a character or numeric value to a date value with the `TO_DATE` function. Oracle automatically converts character values that are in the default date format into date values when they are used in date expressions. The default date format is specified by the initialization parameter `NLS_DATE_FORMAT` and is a string such as `'DD-MON-YY'`. This example date format includes a two-digit number for the day of the month, an abbreviation of the month name, and the last two digits of the year.

If you specify a date value without a time component, the default time is 12:00:00 AM (midnight). If you specify a time value without a date, the default date is the first day of the current month.

The date function `SYSDATE` returns the current date and time. For information on the `SYSDATE` and `TO_DATE` functions and the default date format, see "[Date Format Models](#)" on page 2-45 and [Chapter 4, "Functions"](#).

### Date Arithmetic

You can add and subtract number constants as well as other dates from dates. Oracle interprets number constants in arithmetic date expressions as numbers of days. For example, `SYSDATE + 1` is tomorrow. `SYSDATE - 7` is one week ago. `SYSDATE + (10/1440)` is ten minutes from now. Subtracting the `HIREDATE` column of the `EMP` table from `SYSDATE` returns the number of days since each employee was hired. You cannot multiply or divide `DATE` values.

Oracle provides functions for many common date operations. For example, the `ADD_MONTHS` function lets you add or subtract months from a date. The

MONTHS\_BETWEEN function returns the number of months between two dates. The fractional portion of the result represents that portion of a 31-day month.

Because each date contains a time component, most results of date operations include a fraction. This fraction means a portion of one day. For example, 1.5 days is 36 hours.

**See Also:** ["Date Functions"](#) on page 4-5 for more information on date functions.

### Using Julian Dates

A Julian date is the number of days since January 1, 4712 BC. Julian dates allow continuous dating from a common reference. You can use the date format model "J" with date functions TO\_DATE and TO\_CHAR to convert between Oracle DATE values and their Julian equivalents.

**Example** This statement returns the Julian equivalent of January 1, 1997:

```
SELECT TO_CHAR(TO_DATE('01-01-1997', 'MM-DD-YYYY'), 'J')
       FROM DUAL;
```

```
TO_CHAR
-----
2450450
```

For a description of the DUAL table, see ["Selecting from the DUAL Table"](#) on page 5-26.

## RAW and LONG RAW Datatypes

The RAW and LONG RAW datatypes store data that is not to be interpreted (not explicitly converted when moving data between different systems) by Oracle. These datatypes are intended for binary data or byte strings. For example, you can use LONG RAW to store graphics, sound, documents, or arrays of binary data, for which the interpretation is dependent on the use.

---

---

**Note:** Oracle Corporation strongly recommends that you convert LONG RAW columns to binary LOB (BLOB) columns. LOB columns are subject to far fewer restrictions than LONG columns. See ["TO\\_LOB"](#) on page 4-102 for more information.

---

---

RAW is a variable-length datatype like VARCHAR2, except that Net8 (which connects user sessions to the instance) and the Import and Export utilities do not perform character conversion when transmitting RAW or LONG RAW data. In contrast, Net8 and Import/Export automatically convert CHAR, VARCHAR2, and LONG data from the database character set to the user session character set (which you can set with the NLS\_LANGUAGE parameter of the ALTER SESSION statement), if the two character sets are different.

When Oracle automatically converts RAW or LONG RAW data to and from CHAR data, the binary data is represented in hexadecimal form, with one hexadecimal character representing every four bits of RAW data. For example, one byte of RAW data with bits 11001011 is displayed and entered as 'CB'.

## Large Object (LOB) Datatypes

The built-in LOB datatypes BLOB, CLOB, and NCLOB (stored internally), and the BFILE (stored externally), can store large and unstructured data such as text, image, video, and spatial data up to 4 gigabytes in size.

When creating a table, you can optionally specify different tablespace and storage characteristics for LOB columns or LOB object attributes from those specified for the table.

LOB columns contain LOB locators that can refer to out-of-line or in-line LOB values. Selecting a LOB from a table actually returns the LOB's locator and not the entire LOB value. The DBMS\_LOB package and Oracle Call Interface (OCI) operations on LOBs are performed through these locators.

LOBs are similar to LONG and LONG RAW types, but differ in the following ways:

- LOBs can be attributes of a user-defined datatype (object).
- The LOB locator is stored in the table column, either with or without the actual LOB value. BLOB, NCLOB, and CLOB values can be stored in separate tablespaces. BFILE data is stored in an external file on the server.
- When you access a LOB column, the locator is returned.
- A LOB can be up to 4 gigabytes in size. BFILE maximum size is operating system dependent, but cannot exceed 4 gigabytes.
- LOBs permit efficient, random, piece-wise access to and manipulation of data.
- You can define more than one LOB column in a table.
- With the exception of NCLOB, you can define one or more LOB attributes in an object.

- You can declare LOB bind variables.
- You can select LOB columns and LOB attributes.
- You can insert a new row or update an existing row that contains one or more LOB columns and/or an object with one or more LOB attributes. (You can set the internal LOB value to NULL, empty, or replace the entire LOB with data. You can set the BFILE to NULL or make it point to a different file.)
- You can update a LOB row/column intersection or a LOB attribute with another LOB row/column intersection or LOB attribute.
- You can delete a row containing a LOB column or LOB attribute and thereby also delete the LOB value. Note that for BFILES, the actual operating system file is not deleted.

You can access and populate rows of an internal LOB column (a LOB column stored in the database) simply by issuing an `INSERT` or `UPDATE` statement. However, to access and populate a LOB attribute that is part of an object type, you must first initialize the LOB attribute using the `EMPTY_CLOB` or `EMPTY_BLOB` function. (See "[EMPTY\\_\[B | C\]LOB](#)" on page 4-33.) You can then select the empty LOB attribute and populate it using the `DBMS_LOB` package or some other appropriate interface.

The following example creates a table with LOB columns. (It assumes the existence of tablespace `RESUMES`).

```
CREATE TABLE person_table (name CHAR(40),
                           resume CLOB,
                           picture BLOB)
LOB (resume) STORE AS
    ( TABLESPACE resumes
      STORAGE (INITIAL 5M NEXT 5M) );
```

**See Also:**

- *Oracle8i Supplied PL/SQL Packages Reference* and *Oracle Call Interface Programmer's Guide* for more information about these interfaces and LOBs.
- *Oracle8i Application Developer's Guide - Large Objects (LOBs)* for information on creating temporary LOBs and on LOB restrictions.
- "[TO\\_LOB](#)" on page 4-102 for more information on converting `LONG` columns to LOB columns.



## **BFILE Datatype**

The `BFILE` datatype enables access to binary file LOBs that are stored in file systems outside the Oracle database. A `BFILE` column or attribute stores a `BFILE` locator, which serves as a pointer to a binary file on the server's file system. The locator maintains the directory alias and the filename. See "[CREATE DIRECTORY](#)" on page 7-282.

Binary file LOBs do not participate in transactions and are not recoverable. Rather, the underlying operating system provides file integrity and durability. The maximum file size supported is 4 gigabytes.

The database administrator must ensure that the file exists and that Oracle processes have operating system read permissions on the file.

The `BFILE` datatype allows read-only support of large binary files. You cannot modify or replicate such a file. Oracle provides APIs to access file data. The primary interfaces that you use to access file data are the `DBMS_LOB` package and the OCI.

**See Also:** *Oracle8i Application Developer's Guide - Large Objects (LOBs)* and *Oracle Call Interface Programmer's Guide* for more information about LOBs.

## **BLOB Datatype**

The `BLOB` datatype stores unstructured binary large objects. `BLOBs` can be thought of as bitstreams with no character set semantics. `BLOBs` can store up to 4 gigabytes of binary data.

`BLOBs` have full transactional support. Changes made through SQL, the `DBMS_LOB` package, or the OCI participate fully in the transaction. `BLOB` value manipulations can be committed and rolled back. Note, however, that you cannot save a `BLOB` locator in a PL/SQL or OCI variable in one transaction and then use it in another transaction or session.

## **CLOB Datatype**

The `CLOB` datatype stores single-byte character data. Both fixed-width and variable-width character sets are supported, and both use the `CHAR` database character set. `CLOBs` can store up to 4 gigabytes of character data.

`CLOBs` have full transactional support. Changes made through SQL, the `DBMS_LOB` package, or the OCI participate fully in the transaction. `CLOB` value manipulations can be committed and rolled back. Note, however, that you cannot save a `CLOB`

locator in a PL/SQL or OCI variable in one transaction and then use it in another transaction or session.

### **NCLOB Datatype**

The `NCLOB` datatype stores multibyte national character set character (`NCHAR`) data. Both fixed-width and variable-width character sets are supported. `NCLOBs` can store up to 4 gigabytes of character text data.

`NCLOBs` have full transactional support. Changes made through SQL, the `DBMS_LOB` package, or the OCI participate fully in the transaction. `NCLOB` value manipulations can be committed and rolled back. Note, however, that you cannot save an `NCLOB` locator in a PL/SQL or OCI variable in one transaction and then use it in another transaction or session.

### **ROWID Datatype**

Each row in the database has an address. You can examine a row's address by querying the pseudocolumn `ROWID`. Values of this pseudocolumn are hexadecimal strings representing the address of each row. These strings have the datatype `ROWID`. You can also create tables and clusters that contain actual columns having the `ROWID` datatype. Oracle does not guarantee that the values of such columns are valid rowids.

**See Also:** ["Pseudocolumns"](#) on page 2-56 for more information on the `ROWID` pseudocolumn.

### **Restricted Rowids**

Beginning with Oracle8, Oracle SQL incorporated an extended format for rowids to efficiently support partitioned tables and indexes and tablespace-relative data block addresses (DBAs) without ambiguity.

Character values representing rowids in Oracle7 and earlier releases are called **restricted** rowids. Their format is as follows:

```
block.row.file
```

where:

*block* is a hexadecimal string identifying the data block of the datafile containing the row. The length of this string depends on your operating system.

---

<i>row</i>	is a four-digit hexadecimal string identifying the row in the data block. The first row of the block has a digit of 0.
<i>file</i>	is a hexadecimal string identifying the database file containing the row. The first datafile has the number 1. The length of this string depends on your operating system.

### Extended Rowids

The **extended** ROWID datatype stored in a user column includes the data in the restricted rowid plus a **data object number**. The data object number is an identification number assigned to every database segment. You can retrieve the data object number from data dictionary views USER\_OBJECTS, DBA\_OBJECTS, and ALL\_OBJECTS. Objects that share the same segment (clustered tables in the same cluster, for example) have the same object number.

Extended rowids are not available directly. You can use a supplied package, DBMS\_ROWID, to interpret extended rowid contents. The package functions extract and provide information that would be available directly from a restricted rowid, as well as information specific to extended rowids. For information on the functions available with the DBMS\_ROWID package and how to use them, see *Oracle8i Supplied PL/SQL Packages Reference*.

### Compatibility and Migration

The restricted form of a rowid is still supported in Oracle8i for backward compatibility, but all tables return rowids in the extended format. For information regarding compatibility and migration issues, see *Oracle8i Migration*.

## UROWID Datatype

Each row in a database has an address (as discussed in "[ROWID Datatype](#)" on page 2-26). However, the rows of some tables have addresses that are not physical or permanent, or were not generated by Oracle. For example, the row addresses of index-organized tables are stored in index leaves, which can move. Rowids of foreign tables (such as DB2 tables accessed through a gateway) are not standard Oracle rowids.

Oracle uses "universal rowids" (**urowids**) to store the addresses of index-organized and foreign tables. Index-organized tables have logical urowids and foreign tables have foreign urowids. Both types of urowid are stored in the ROWID pseudocolumn (as are the physical rowids of heap-organized tables).

Oracle creates logical rowids based on a table's primary key. The logical rowids do not change as long as the primary key does not change. The ROWID pseudocolumn of an index-organized table has a datatype of UROWID. You can access this pseudocolumn as you would the ROWID pseudocolumn of a heap-organized (that is, using the SELECT ROWID statement). If you wish to store the rowids of an index-organized table, you can define a column of type UROWID for the table and retrieve the value of the ROWID pseudocolumn into that column.

---

---

**Note:** Heap-organized tables have physical rowids. Oracle Corporation does not recommend that you specify a column of datatype UROWID for a heap-organized table.

---

---

**See Also:** *Oracle8i Concepts* and *Oracle8i Designing and Tuning for Performance* for more information on the UROWID datatype and how Oracle generates and manipulates universal rowids

## ANSI, DB2, and SQL/DS Datatypes

SQL statements that create tables and clusters can also use ANSI datatypes and datatypes from IBM's products SQL/DS and DB2. Oracle recognizes the ANSI or IBM datatype name and records it as the name of the datatype of the column, and then stores the column's data in an Oracle datatype based on the conversions shown in [Table 2-2](#) and [Table 2-3](#).

**Table 2-2 ANSI Datatypes Converted to Oracle Datatypes**

ANSI SQL Datatype	Oracle Datatype
CHARACTER (n)	CHAR (n)
CHAR (n)	
CHARACTER VARYING (n)	VARCHAR (n)
CHAR VARYING (n)	
NATIONAL CHARACTER (n)	NCHAR (n)
NATIONAL CHAR (n)	
NCHAR (n)	

**Table 2–2 ANSI Datatypes Converted to Oracle Datatypes**

NATIONAL CHARACTER VARYING(n)	NVARCHAR2(n)
NATIONAL CHAR VARYING(n)	
NCHAR VARYING(n)	
NUMERIC(p, s)	NUMBER(p, s)
DECIMAL(p, s) <sup>a</sup>	
INTEGER	NUMBER(38)
INT	
SMALLINT	
FLOAT(b) <sup>b</sup>	NUMBER
DOUBLE PRECISION <sup>c</sup>	
REAL <sup>d</sup>	

<sup>a</sup>The NUMERIC and DECIMAL datatypes can specify only fixed-point numbers. For these datatypes, *s* defaults to 0.

<sup>b</sup>The FLOAT datatype is a floating-point number with a binary precision *b*. The default precision for this datatype is 126 binary, or 38 decimal.

<sup>c</sup>The DOUBLE PRECISION datatype is a floating-point number with binary precision 126.

<sup>d</sup>The REAL datatype is a floating-point number with a binary precision of 63, or 18 decimal.

**Table 2–3 SQL/DS and DB2 Datatypes Converted to Oracle Datatypes**

SQL/DS or DB2 Datatype	Oracle Datatype
CHARACTER(n)	CHAR(n)
VARCHAR(n)	VARCHAR(n)
LONG VARCHAR(n)	LONG
DECIMAL(p, s) <sup>a</sup>	NUMBER(p, s)
INTEGER	NUMBER(38)
SMALLINT	
FLOAT(b) <sup>b</sup>	NUMBER

**Table 2–3 SQL/DS and DB2 Datatypes Converted to Oracle Datatypes**

<sup>a</sup>The DECIMAL datatype can specify only fixed-point numbers. For this datatype, *s* defaults to 0.

<sup>b</sup>The FLOAT datatype is a floating-point number with a binary precision *b*. The default precision for this datatype is 126 binary, or 38 decimal.

---

Do not define columns with these SQL/DS and DB2 datatypes, because they have no corresponding Oracle datatype:

- GRAPHIC
- LONG VARGRAPHIC
- VARGRAPHIC
- TIME
- TIMESTAMP

Note that data of type TIME and TIMESTAMP can also be expressed as Oracle DATE data.

## User-Defined Type Categories

User-defined datatypes use Oracle built-in datatypes and other user-defined datatypes as the building blocks of types that model the structure and behavior of data in applications. For information about Oracle built-in datatypes, see *Oracle8i Concepts*. For information about creating user-defined types, see "CREATE TYPE" on page 7-437 and the "CREATE TYPE BODY" on page 7-447. For information about using user-defined types, see *Oracle8i Application Developer's Guide - Fundamentals*.

The sections that follow describe the various categories of user-defined types.

### Object Types

Object types are abstractions of the real-world entities, such as purchase orders, that application programs deal with. An object type is a schema object with three kinds of components:

- A *name*, which identifies the object type uniquely within that schema.
- *Attributes*, which are built-in types or other user-defined types. Attributes model the structure of the real-world entity.
- *Methods*, which are functions or procedures written in PL/SQL and stored in the database, or written in a language like C or Java and stored externally.

Methods implement operations the application can perform on the real-world entity.

## REFs

An **object identifier** (OID) uniquely identifies an object and enables you to reference the object from other objects or from relational tables. A datatype category called REF represents such references. A REF is a container for an object identifier. REFs are pointers to objects.

When a REF value points to a nonexistent object, the REF is said to be "dangling". A dangling REF is different from a null REF. To determine whether a REF is dangling or not, use the predicate `IS [NOT] DANGLING`. For example, given table DEPT with column MGR whose type is a REF to type EMP\_T, which has an attribute NAME:

```
SELECT t.mgr.name
       FROM dept t
       WHERE t.mgr IS NOT DANGLING;
```

## Varrays

An array is an ordered set of data elements. All elements of a given array are of the same datatype. Each element has an **index**, which is a number corresponding to the element's position in the array.

The number of elements in an array is the size of the array. Oracle arrays are of variable size, which is why they are called **varrays**. You must specify a maximum size when you declare the array.

When you declare a varray, it does not allocate space. It defines a type, which you can use as:

- The datatype of a column of a relational table
- An object type attribute
- A PL/SQL variable, parameter, or function return type

Oracle normally stores an array object either in line (that is, as part of the row data) or out of line (in a LOB), depending on its size. However, if you specify separate storage characteristics for a varray, Oracle will store it out of line, regardless of its size (see the *[varray\\_storage\\_clause](#)* of "CREATE TABLE" on page 4-381).

## Nested Tables

A nested table type models an unordered set of elements. The elements may be built-in types or user-defined types. You can view a nested table as a single-column

table or, if the nested table is an object type, as a multicolumn table, with a column for each attribute of the object type.

A nested table definition does not allocate space. It defines a type, which you can use to declare:

- Columns of a relational table
- Object type attributes
- PL/SQL variables, parameters, and function return values

When a nested table appears as the type of a column in a relational table or as an attribute of the underlying object type of an object table, Oracle stores all of the nested table data in a single table, which it associates with the enclosing relational or object table.

## Datatype Comparison Rules

This section describes how Oracle compares values of each datatype.

### Number Values

A larger value is considered greater than a smaller one. All negative numbers are less than zero and all positive numbers. Thus, -1 is less than 100; -100 is less than -1.

### Date Values

A later date is considered greater than an earlier one. For example, the date equivalent of '29-MAR-1997' is less than that of '05-JAN-1998' and '05-JAN-1998 1:35pm' is greater than '05-JAN-1998 10:09am'.

### Character String Values

Character values are compared using one of these comparison rules:

- blank-padded comparison semantics
- nonpadded comparison semantics

The following sections explain these comparison semantics. The results of comparing two character values using different comparison semantics may vary. The table below shows the results of comparing five pairs of character values using each comparison semantic. Usually, the results of blank-padded and nonpadded comparisons are the same. The last comparison in the table illustrates the differences between the blank-padded and nonpadded comparison semantics.



Blank-Padded	Nonpadded
'ab' > 'aa'	'ab' > 'aa'
'ab' > 'a '	'ab' > 'a '
'ab' > 'a'	'ab' > 'a'
'ab' = 'ab'	'ab' = 'ab'
'a ' = 'a'	'a ' > 'a'

**Blank-Padded Comparison Semantics** If the two values have different lengths, Oracle first adds blanks to the end of the shorter one so their lengths are equal. Oracle then compares the values character by character up to the first character that differs. The value with the greater character in the first differing position is considered greater. If two values have no differing characters, then they are considered equal. This rule means that two values are equal if they differ only in the number of trailing blanks. Oracle uses blank-padded comparison semantics only when both values in the comparison are either expressions of datatype `CHAR`, `NCHAR`, text literals, or values returned by the `USER` function.

**Nonpadded Comparison Semantics** Oracle compares two values character by character up to the first character that differs. The value with the greater character in that position is considered greater. If two values of different length are identical up to the end of the shorter one, the longer value is considered greater. If two values of equal length have no differing characters, then the values are considered equal. Oracle uses nonpadded comparison semantics whenever one or both values in the comparison have the datatype `VARCHAR2` or `NVARCHAR2`.

### Single Characters

Oracle compares single characters according to their numeric values in the database character set. One character is greater than another if it has a greater numeric value than the other in the character set. Oracle considers blanks to be less than any character, which is true in most character sets.

These are some common character sets:

- 7-bit ASCII (American Standard Code for Information Interchange)
- EBCDIC Code (Extended Binary Coded Decimal Interchange Code) Page 500
- ISO 8859/1 (International Standards Organization)
- JEUC Japan Extended UNIX

Portions of the ASCII and EBCDIC character sets appear in [Table 2-4](#) and [Table 2-5](#). Note that uppercase and lowercase letters are not equivalent. Also, note that the numeric values for the characters of a character set may not match the linguistic sequence for a particular language.

**Table 2-4** *ASCII Character Set*

Symbol	Decimal value	Symbol	Decimal value
blank	32	;	59
!	33	<	60
"	34	=	61
#	35	>	62
\$	36	?	63
%	37	@	64
&	38	A-Z	65-90
'	39	[	91
(	40	\	92
)	41	]	93
*	42	^	94
+	43	_	95
,	44	`	96
-	45	a-z	97-122
.	46	{	123
/	47		124
0-9	48-57	}	125
:	58	~	126

**Table 2–5 EBCDIC Character Set**

Symbol	Decimal value	Symbol	Decimal value
blank	64	%	108
¢	74	—	109
.	75	>	110
<	76	?	111
(	77	:	122
+	78	#	123
	79	@	124
&	80	'	125
!	90	=	126
\$	91	"	127
*	92	a-i	129-137
)	93	j-r	145-153
;	94	s-z	162-169
ÿ	95	A-I	193-201
-	96	J-R	209-217
/	97	S-Z	226-233

### Object Values

Object values are compared using one of two comparison functions: `MAP` and `ORDER`. Both functions compare object type instances, but they are quite different from one another. These functions must be specified as part of the object type.

For a description of `MAP` and `ORDER` methods and the values they return, see "[CREATE TYPE](#)" on page 7-437. See also *Oracle8i Application Developer's Guide - Fundamentals* for more information.

### Varrays and Nested Tables

You cannot compare varrays and nested tables in Oracle8i.

## Data Conversion

Generally an expression cannot contain values of different datatypes. For example, an expression cannot multiply 5 by 10 and then add 'JAMES'. However, Oracle supports both implicit and explicit conversion of values from one datatype to another.

### Implicit Data Conversion

Oracle automatically converts a value from one datatype to another when such a conversion makes sense. Oracle performs conversions in these cases:

- When an `INSERT` or `UPDATE` statement assigns a value of one datatype to a column of another, Oracle converts the value to the datatype of the column.
- When you use a SQL function or operator with an argument with a datatype other than the one it accepts, Oracle converts the argument to the accepted datatype.
- When you use a comparison operator on values of different datatypes, Oracle converts one of the expressions to the datatype of the other.

**Example 1** The text literal '10' has datatype `CHAR`. Oracle implicitly converts it to the `NUMBER` datatype if it appears in a numeric expression as in the following statement:

```
SELECT sal + '10'  
       FROM emp;
```

**Example 2** When a condition compares a character value and a `NUMBER` value, Oracle implicitly converts the character value to a `NUMBER` value, rather than converting the `NUMBER` value to a character value. In the following statement, Oracle implicitly converts '7936' to 7936:

```
SELECT ename  
       FROM emp  
       WHERE empno = '7936';
```

**Example 3** In the following statement, Oracle implicitly converts '12-MAR-1993' to a `DATE` value using the default date format 'DD-MON-YYYY':

```
SELECT ename  
       FROM emp  
       WHERE hiredate = '12-MAR-1993';
```

**Example 4** In the following statement, Oracle implicitly converts the text literal 'AAAAZ8AABAAABv1AAA' to a rowid value:

```
SELECT ename
       FROM emp
       WHERE ROWID = 'AAAAZ8AABAAABv1AAA';
```

### Explicit Data Conversion

You can also explicitly specify datatype conversions using SQL conversion functions. [Table 2-6](#) shows SQL functions that explicitly convert a value from one datatype to another.

**Table 2-6 SQL Functions for Datatype Conversion**

TO / FROM	CHAR	NUMBER	DATE	RAW	ROWID	LONG/ LONG RAW	LOB
<b>CHAR</b>	—	TO_NUMBER	TO_DATE	HEXTORAW	CHARTO- ROWID		
<b>NUMBER</b>	TO_CHAR	—	TO_DATE (number, ' J')				
<b>DATE</b>	TO_CHAR	TO_CHAR (date, 'J' )	—				
<b>RAW</b>	RAWTOHEX			—			
<b>ROWID</b>	ROWID- TOCHAR				—		
<b>LONG / LONG RAW</b>						—	TO_LOB
<b>LOB</b>							—

For information on these functions, see "[Conversion Functions](#)" on page 4-5.

---

---

**Note:** You cannot specify `LONG` and `LONG RAW` values in cases in which Oracle can perform implicit datatype conversion. For example, `LONG` and `LONG RAW` values cannot appear in expressions with functions or operators. For information on the limitations on `LONG` and `LONG RAW` datatypes, see "[LONG Datatype](#)" on page 2-19.

---

---

### Implicit vs. Explicit Data Conversion

Oracle recommends that you specify explicit conversions rather than rely on implicit or automatic conversions for these reasons:

- SQL statements are easier to understand when you use explicit datatype conversion functions.
- Automatic datatype conversion can have a negative impact on performance, especially if the datatype of a column value is converted to that of a constant rather than the other way around.
- Implicit conversion depends on the context in which it occurs and may not work the same way in every case.
- Algorithms for implicit conversion are subject to change across software releases and among Oracle products. Behavior of explicit conversions is more predictable.

## Format Models

A **format model** is a character literal that describes the format of `DATE` or `NUMBER` data stored in a character string. You can use a format model as an argument of the `TO_CHAR` and `TO_DATE` functions:

- To specify the format for Oracle to use to return a value from the database
- To specify the format for a value you have specified for Oracle to store in the database

See "[TO\\_CHAR \(date conversion\)](#)" on page 4-99, "[TO\\_CHAR \(number conversion\)](#)" on page 4-100, and "[TO\\_DATE](#)" on page 4-101. Note that a format model does not change the internal representation of the value in the database.

This section describes how to use:

- Number format models
- Date format models

- Format model modifiers

## Changing the Return Format

You can use a format model to specify the format for Oracle to use to return values from the database to you.

**Example 1** The following statement selects the commission values of the employees in Department 30 and uses the `TO_CHAR` function to convert these commissions into character values with the format specified by the number format model `'$9,990.99'`:

```
SELECT ename employee, TO_CHAR(comm, '$9,990.99') commission
       FROM emp
       WHERE deptno = 30;
```

EMPLOYEE	COMMISSION
ALLEN	\$300.00
WARD	\$500.00
MARTIN	\$1,400.00
BLAKE	
TURNER	\$0.00
JAMES	

Because of this format model, Oracle returns commissions with leading dollar signs, commas every three digits, and two decimal places. Note that `TO_CHAR` returns null for all employees with null in the `COMM` column.

**Example 2** The following statement selects the date on which each employee from Department 20 was hired and uses the `TO_CHAR` function to convert these dates to character strings with the format specified by the date format model `'fmMonth DD, YYYY'`:

```
SELECT ename, TO_CHAR(Hiredate, 'fmMonth DD, YYYY') hiredate
       FROM emp
       WHERE deptno = 20;
```

ENAME	HIREDATE
SMITH	December 17, 1980
JONES	April 2, 1981
SCOTT	April 19, 1987

```
ADAMS      May 23, 1987
FORD       December 3, 1981
LEWIS      October 23, 1997
```

With this format model, Oracle returns the hire dates (as specified by "fm" and discussed in "[Format Model Modifiers](#)" on page 2-51) without blank padding, two digits for the day, and the century included in the year.

## Supplying the Correct Format

You can use format models to specify the format of a value that you are converting from one datatype to another datatype required for a column. When you insert or update a column value, the datatype of the value that you specify must correspond to the column's datatype. For example, a value that you insert into a DATE column must be a value of the DATE datatype or a character string in the default date format (Oracle implicitly converts character strings in the default date format to the DATE datatype). If the value is in another format, you must use the TO\_DATE function to convert the value to the DATE datatype. You must also use a format model to specify the format of the character string.

**Example** The following statement updates BAKER's hire date using the TO\_DATE function with the format mask 'YYYY MM DD' to convert the character string '1998 05 20' to a DATE value:

```
UPDATE emp
   SET hiredate = TO_DATE('1998 05 20', 'YYYY MM DD')
   WHERE ename = 'BLAKE';
```

## Number Format Models

You can use number format models:

- In the TO\_CHAR function to translate a value of NUMBER datatype to VARCHAR2 datatype
- In the TO\_NUMBER function to translate a value of CHAR or VARCHAR2 datatype to NUMBER datatype

All number format models cause the number to be rounded to the specified number of significant digits. If a value has more significant digits to the left of the decimal place than are specified in the format, pound signs (#) replace the value. If a positive value is extremely large and cannot be represented in the specified format, then the infinity sign (~) replaces the value. Likewise, if a negative value is extremely small and cannot be represented by the specified format, then the negative infinity sign



replaces the value (--). This event typically occurs when you are using `TO_CHAR` with a restrictive number format string, causing a rounding operation.

## Number Format Elements

A number format model is composed of one or more number format elements. [Table 2-7](#) lists the elements of a number format model. Examples are shown in [Table 2-8](#).

Negative return values automatically contain a leading negative sign and positive values automatically contain a leading space unless the format model contains the `MI`, `S`, or `PR` format element.

**Table 2-7** *Number Format Elements*

Element	Example	Description
, (comma)	9,999	Returns a comma in the specified position. You can specify multiple commas in a number format model.  <b>Restrictions:</b> <ul style="list-style-type: none"> <li>■ A comma element cannot begin a number format model.</li> <li>■ A comma cannot appear to the right of a decimal character or period in a number format model.</li> </ul>
. (period)	99.99	Returns a decimal point, which is a period (.) in the specified position.  <b>Restriction:</b> You can specify only one period in a number format model.
\$	\$9999	Returns value with a leading dollar sign.
0	0999 9990	Returns leading zeros. Returns trailing zeros.
9	9999	Returns value with the specified number of digits with a leading space if positive or with a leading minus if negative. Leading zeros are blank, except for a zero value, which returns a zero for the integer part of the fixed-point number.
B	B9999	Returns blanks for the integer part of a fixed-point number when the integer part is zero (regardless of "0"s in the format model).
C	C999	Returns in the specified position the ISO currency symbol (the current value of the <code>NLS_ISO_CURRENCY</code> parameter).

**Table 2–7 (Cont.) Number Format Elements**

Element	Example	Description
D	99D99	Returns in the specified position the decimal character, which is the current value of the <code>NLS_NUMERIC_CHARACTER</code> parameter. The default is a period (.).  <b>Restriction:</b> You can specify only one decimal character in a number format model.
EEEE	9.9EEEE	Returns a value using in scientific notation.
FM	FM90.9	Returns a value with no leading or trailing blanks.
G	9G999	Returns in the specified position the group separator (the current value of the <code>NLS_NUMERIC_CHARACTER</code> parameter). You can specify multiple group separators in a number format model.  <b>Restriction:</b> A group separator cannot appear to the right of a decimal character or period in a number format model.
L	L999	Returns in the specified position the local currency symbol (the current value of the <code>NLS_CURRENCY</code> parameter).
MI	9999MI	Returns negative value with a trailing minus sign (-). Returns positive value with a trailing blank.  <b>Restriction:</b> The MI format element can appear only in the last position of a number format model.
PR	9999PR	Returns negative value in <angle brackets>. Returns positive value with a leading and trailing blank.  <b>Restriction:</b> The PR format element can appear only in the last position of a number format model.
RN	RN	Returns a value as Roman numerals in uppercase.
rn	rn	Returns a value as Roman numerals in lowercase. Value can be an integer between 1 and 3999.
S	S9999	Returns negative value with a leading minus sign (-). Returns positive value with a leading plus sign (+).
	9999S	Returns negative value with a trailing minus sign (-). Returns positive value with a trailing plus sign (+).  <b>Restriction:</b> The S format element can appear only in the first or last position of a number format model.

**Table 2–7 (Cont.) Number Format Elements**

Element	Example	Description
TM	TM	<p>"Text minimum". Returns (in decimal output) the smallest number of characters possible. This element is case-insensitive.</p> <p>The default is TM9, which returns the number in fixed notation unless the output exceeds 64 characters. If output exceeds 64 characters, Oracle automatically returns the number in scientific notation.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot precede this element with any other element.</li> <li>■ You can follow this element only with 9 or E (only one) or e (only one).</li> </ul>
U	U9999	Returns in the specified position the "Euro" (or other) dual currency symbol (the current value of the NLS_DUAL_CURRENCY parameter).
V	999V99	Returns a value multiplied by $10^n$ (and if necessary, round it up), where $n$ is the number of 9's after the "V".
X	XXXX xxxx	<p>Returns the hexadecimal value of the specified number of digits. If the specified number is not an integer, Oracle rounds it to an integer.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ This element accepts only positive values or 0. Negative values return an error.</li> <li>■ You can precede this element only with 0 (which returns leading zeroes) or FM. Any other elements return an error. If you specify neither 0 nor FM with X, the return always has 1 leading blank.</li> </ul>

The values of some formats are determined by the value of initialization parameters. For such formats, you can specify the characters returned by these format elements implicitly using the initialization parameter NLS\_TERRITORY. For information on these parameters, see *Oracle8i Reference* and *Oracle8i National Language Support Guide*.

You can change the default date format for your session with the ALTER SESSION statement. For information on changing the settings of these parameters, see "[ALTER SESSION](#)" on page 7-83.

**Example** Table 2-8 shows the results of the following query for different values of *number* and *'fmt'*:

```
SELECT TO_CHAR(number, 'fmt')
       FROM DUAL;
```

**Table 2-8 Results of Example Number Conversions**

number	'fmt'	Result
-1234567890	9999999999S	'1234567890-'
0	99.99	' .00'
+0.1	99.99	' .10'
-0.2	99.99	' -.20'
0	90.99	' 0.00'
+0.1	90.99	' 0.10'
-0.2	90.99	' -0.20'
0	9999	' 0'
1	9999	' 1'
0	B9999	' '
1	B9999	' 1'
0	B90.99	' '
+123.456	999.999	' 123.456'
-123.456	999.999	' -123.456'
+123.456	FM999.009	'123.456'
+123.456	9.9E0000	' 1.2E+02'
+1E+123	9.9E0000	' 1.0E+123'
+123.456	FM9.9E0000	'1.23E+02'
+123.45	FM999.009	'123.45'
+123.0	FM999.009	'123.00'
+123.45	L999.99	' \$123.45'
+123.45	FML99.99	'\$123.45'
+1234567890	9999999999S	'1234567890+'

## Date Format Models

You can use date format models:

- In the `TO_CHAR` function to translate a `DATE` value that is in a format other than the default date format
- In the `TO_DATE` function to translate a character value that is in a format other than the default date format

### Default Date Format

The default date format is specified either explicitly with the initialization parameter `NLS_DATE_FORMAT` or implicitly with the initialization parameter `NLS_TERRITORY`. For information on these parameters, see *Oracle8i Reference*.

You can change the default date format for your session with the `ALTER SESSION` statement. For information, see "[ALTER SESSION](#)" on page 7-83.

### Maximum Length

The total length of a date format model cannot exceed 22 characters.

### Date Format Elements

A date format model is composed of one or more date format elements as listed in [Table 2-9](#).

- For input format models, format items cannot appear twice, and format items that represent similar information cannot be combined. For example, you cannot use 'SYYYY' and 'BC' in the same format string.
- Some of the date format elements cannot be used in the `TO_DATE` function, as noted in [Table 2-9](#).

**Capitalization of Date Format Elements** Capitalization in a spelled-out word, abbreviation, or Roman numeral follows capitalization in the corresponding format element. For example, the date format model 'DAY' produces capitalized words like 'MONDAY'; 'Day' produces 'Monday'; and 'day' produces 'monday'.

**Punctuation and Character Literals in Date Format Models** You can also include these characters in a date format model:

- punctuation such as hyphens, slashes, commas, periods, and colons
- character literals, enclosed in double quotation marks

These characters appear in the return value in the same location as they appear in the format model.

**Table 2–9 Datetime Format Elements**

Element	Specify in TO_DATE?	Meaning
- / , . ; : 'text'	Yes	Punctuation and quoted text is reproduced in the result.
AD A.D.	Yes	AD indicator with or without periods.
AM A.M.	Yes	Meridian indicator with or without periods.
BC B.C.	Yes	BC indicator with or without periods.
CC SCC	No	The first two digits of the century of a four-digit year, for example, '19' from '1900' and '20' from '2001'. "S" prefixes BC dates with "-".
D	Yes	Day of week (1-7). This element is used <i>only</i> to validate a date specified in the TO_DATE function.
DAY	Yes	Name of day, padded with blanks to length of 9 characters. This element is used <i>only</i> to validate a date specified in the TO_DATE function.
DD	Yes	Day of month (1-31).
DDD	Yes	Day of year (1-366).
DY	Yes	Abbreviated name of day. This element is used <i>only</i> to validate a date specified in the TO_DATE function.
E	Yes	Abbreviated era name (Japanese Imperial, ROC Official, and Thai Buddha calendars).
EE	Yes	Full era name (Japanese Imperial, ROC Official, and Thai Buddha calendars).
HH	Yes	Hour of day (1-12).
HH12	Yes	Hour of day (1-12).

**Table 2–9 (Cont.) Datetime Format Elements**

<b>Element</b>	<b>Specify in TO_DATE?</b>	<b>Meaning</b>
HH24	Yes	Hour of day (0-23).
IW	No	Week of year (1-52 or 1-53) based on the ISO standard.
IYY IY I	No	Last 3, 2, or 1 digit(s) of ISO year.
IYYY	No	4-digit year based on the ISO standard.
J	Yes	Julian day; the number of days since January 1, 4712 BC. Number specified with 'J' must be integers.
MI	Yes	Minute (0-59).
MM	Yes	Two-digit numeric abbreviation of month (01-12; JAN = 01)
MON	Yes	Abbreviated name of month.
MONTH	Yes	Name of month, padded with blanks to length of 9 characters.
PM P . M .	No	Meridian indicator with or without periods.
Q	No	Quarter of year (1, 2, 3, 4; JAN-MAR = 1)
RM	Yes	Roman numeral month (I-XII; JAN = I).
RR	Yes	Given a year with 2 digits: <ul style="list-style-type: none"> <li>■ If the year is &lt;50 and the last 2 digits of the current year are &gt;=50, the first 2 digits of the returned year are 1 greater than the first two digits of the current year.</li> <li>■ If the year is &gt;=50 and the last 2 digits of the current year are &lt;50, the first 2 digits of the returned year are the same as the first 2 digits of the current year.</li> </ul>
RRRR	Yes	Round year. Accepts either 4-digit or 2-digit input. If 2-digit, provides the same return as RR. If you don't want this functionality, enter the 4-digit year.
SS	Yes	Second (0-59).
SSSS	Yes	Seconds past midnight (0-86399).

**Table 2–9 (Cont.) Datetime Format Elements**

Element	Specify in TO_DATE?	Meaning
WW	No	Week of year (1-53) where week 1 starts on the first day of the year and continues to the seventh day of the year.
W	No	Week of month (1-5) where week 1 starts on the first day of the month and ends on the seventh.
Y, YYY	Yes	Year with comma in this position.
YEAR SYEAR	No	Year, spelled out. "S" prefixes BC dates with "-".
YYYY SYYYY	Yes	4-digit year. "S" prefixes BC dates with "-".
YYY YY Y	Yes	Last 3, 2, or 1 digit(s) of year.

Oracle returns an error if an alphanumeric character is found in the date string where punctuation character is found in the format string. For example:

```
TO_CHAR (TO_DATE('0297', 'MM/YY'), 'MM/YY')
```

returns an error.

### Date Format Elements and National Language Support

The functionality of some date format elements depends on the country and language in which you are using Oracle. For example, these date format elements return spelled values:

- MONTH
- MON
- DAY
- DY
- BC or AD or B.C. or A.D.
- AM or PM or A.M or P.M.

The language in which these values are returned is specified either explicitly with the initialization parameter `NLS_DATE_LANGUAGE` or implicitly with the



initialization parameter `NLS_LANGUAGE`. The values returned by the `YEAR` and `SYEAR` date format elements are always in English.

The date format element `D` returns the number of the day of the week (1-7). The day of the week that is numbered 1 is specified implicitly by the initialization parameter `NLS_TERRITORY`.

For information on national language support initialization parameters, see *Oracle8i Reference* and *Oracle8i National Language Support Guide*.

## ISO Standard Date Format Elements

Oracle calculates the values returned by the date format elements `IYYY`, `IYY`, `IY`, `I`, and `IW` according to the ISO standard. For information on the differences between these values and those returned by the date format elements `YYYY`, `YYY`, `YY`, `Y`, and `WW`, see the discussion of national language support in *Oracle8i National Language Support Guide*.

## The RR Date Format Element

The `RR` date format element is similar to the `YY` date format element, but it provides additional flexibility for storing date values in other centuries. The `RR` date format element allows you to store 21st century dates in the 20th century by specifying only the last two digits of the year. It will also allow you to store 20th century dates in the 21st century in the same way if necessary.

If you use the `TO_DATE` function with the `YY` date format element, the date value returned is always in the current century. If you use the `RR` date format element instead, the century of the return value varies according to the specified two-digit year and the last two digits of the current year. [Table 2-10](#) summarizes the behavior of the `RR` date format element.

**Table 2-10 The RR Date Element Format**

		If the specified two-digit year is	
		0 - 49	50 - 99
If the last two digits of the current year are:	0-49	The return date has the same first 2 digits as the current date.	The first 2 digits of the return date are 1 less than the first 2 digits of the current date.
	50-99	The first 2 digits of the return date are 1 greater than the first 2 digits of the current date.	The return date has the same first 2 digits as the current date.

The following examples demonstrate the behavior of the RR date format element.

**Example 1** Assume these queries are issued between 1950 and 1999:

```
SELECT TO_CHAR(TO_DATE('27-OCT-98', 'DD-MON-RR'), 'YYYY') "Year"
       FROM DUAL;
```

Year

----

1998

```
SELECT TO_CHAR(TO_DATE('27-OCT-17', 'DD-MON-RR'), 'YYYY') "Year"
       FROM DUAL;
```

Year

----

2017

**Example 2** Assume these queries are issued between 2000 and 2049:

```
SELECT TO_CHAR(TO_DATE('27-OCT-98', 'DD-MON-RR'), 'YYYY') "Year"
       FROM DUAL;
```

Year

----

1998

```
SELECT TO_CHAR(TO_DATE('27-OCT-17', 'DD-MON-RR'), 'YYYY') "Year"
       FROM DUAL;
```

Year

----

2017

Note that the queries return the same values regardless of whether they are issued before or after the year 2000. The RR date format element allows you to write SQL statements that will return the same values after the turn of the century.

### Date Format Element Suffixes

[Table 2-11](#) lists suffixes that can be added to date format elements:

**Table 2–11 Date Format Element Suffixes**

Suffix	Meaning	Example Element	Example Value
TH	Ordinal Number	DDTH	4TH
SP	Spelled Number	DDSP	FOUR
SPTH or THSP	Spelled, ordinal number	DDSPTH	FOURTH

**Restrictions:**

- When you add one of these suffixes to a date format element, the return value is always in English.
- Date suffixes are valid only on output. You cannot use them to insert a date into the database.

## Format Model Modifiers

The FM and FX modifiers, used in format models in the `TO_CHAR` function, control blank padding and exact format checking.

A modifier can appear in a format model more than once. In such a case, each subsequent occurrence toggles the effects of the modifier. Its effects are enabled for the portion of the model following its first occurrence, and then disabled for the portion following its second, and then reenabled for the portion following its third, and so on.

**FM** "Fill mode". This modifier suppresses blank padding in the return value of the `TO_CHAR` function:

- In a date format element of a `TO_CHAR` function, this modifier suppresses blanks in subsequent character elements (such as MONTH) and suppresses leading zeroes for subsequent number elements (such as MI) in a date format model. Without FM, the result of a character element is always right padded with blanks to a fixed length, and leading zeroes are always returned for a number element. With FM, because there is no blank padding, the length of the return value may vary.
- In a number format element of a `TO_CHAR` function, this modifier suppresses blanks added to the left of the number, so that the result is left-justified in the output buffer. Without FM, the result is always right-justified in the buffer, resulting in blank-padding to the left of the number.

**FX** "Format exact". This modifier specifies exact matching for the character argument and date format model of a `TO_DATE` function:

- Punctuation and quoted text in the character argument must exactly match (except for case) the corresponding parts of the format model.
- The character argument cannot have extra blanks. Without FX, Oracle ignores extra blanks.
- Numeric data in the character argument must have the same number of digits as the corresponding element in the format model. Without FX, numbers in the character argument can omit leading zeroes.

When FX is enabled, you can disable this check for leading zeroes by using the FM modifier as well.

If any portion of the character argument violates any of these conditions, Oracle returns an error message.

**Example 1** The following statement uses a date format model to return a character expression:

```
SELECT TO_CHAR(SYSDATE, 'fmDDTH')||' of '||TO_CHAR
       (SYSDATE, 'fmMonth')||', '||TO_CHAR(SYSDATE, 'YYYY') "Ides"
FROM DUAL;
```

```
Ides
-----
3RD of April, 1998
```

Note that the statement above also uses the FM modifier. If FM is omitted, the month is blank-padded to nine characters:

```
SELECT TO_CHAR(SYSDATE, 'DDTH')||' of '||
       TO_CHAR(SYSDATE, 'Month')||', '||
       TO_CHAR(SYSDATE, 'YYYY') "Ides"
FROM DUAL;
```

```
Ides
-----
03RD of April      , 1998
```

**Example 2** The following statement places a single quotation mark in the return value by using a date format model that includes two consecutive single quotation marks:

```
SELECT TO_CHAR(SYSDATE, 'fmDay')||''''s Special' "Menu"
FROM DUAL;
```

```
Menu
-----
Tuesday's Special
```

Two consecutive single quotation marks can be used for the same purpose within a character literal in a format model.

**Example 3** Table 2-12 shows whether the following statement meets the matching conditions for different values of `char` and `'fmt'` using `FX` (the table named `TABLE` has a column `DATE_COLUMN` of datatype `DATE`):

```
UPDATE table
   SET date_column = TO_DATE(char, 'fmt');
```

**Table 2-12 Matching Character Data and Format Models with the `FX` Format Model Modifier**

char	'fmt'	Match or Error?
'15/ JAN /1998'	'DD-MON-YYYY'	Match
' 15! JAN % /1998'	'DD-MON-YYYY'	Error
'15/JAN/1998'	'FXDD-MON-YYYY'	Error
'15-JAN-1998'	'FXDD-MON-YYYY'	Match
'1-JAN-1998'	'FXDD-MON-YYYY'	Error
'01-JAN-1998'	'FXDD-MON-YYYY'	Match
'1-JAN-1998'	'FXFMDD-MON-YYYY'	Match

## String-to-Date Conversion Rules

The following additional formatting rules apply when converting string values to date values (*unless you have used the `FX` or `FXFM` modifiers in the format model to control exact format checking*):

- You can omit punctuation included in the format string from the date string if all the digits of the numerical format elements, including leading zeros, are specified. In other words, specify `02` and not `2` for two-digit format elements such as `MM`, `DD`, and `YY`.
- You can omit time fields found at the end of a format string from the date string.

- If a match fails between a date format element and the corresponding characters in the date string, Oracle attempts alternative format elements, as shown in [Table 2-13](#).

**Table 2-13 Oracle Format Matching**

Original Format Element	Additional Format Elements to Try in Place of the Original
'MM'	'MON' and 'MONTH'
'MON'	'MONTH'
'MONTH'	'MON'
'YY'	'YYYY'
'RR'	'RRRR'

## Nulls

If a column in a row has no value, then the column is said to be **null**, or to contain a null. Nulls can appear in columns of any datatype that are not restricted by `NOT NULL` or `PRIMARY KEY` integrity constraints. Use a null when the actual value is not known or when a value would not be meaningful.

Do not use null to represent a value of zero, because they are not equivalent. (Oracle currently treats a character value with a length of zero as null. However, this may not continue to be true in future releases, and Oracle recommends that you do not treat empty strings the same as nulls.) Any arithmetic expression containing a null always evaluates to null. For example, null added to 10 is null. In fact, all operators (except concatenation) return null when given a null operand.

## Nulls in SQL Functions

All scalar functions (except `REPLACE`, `NVL`, and `CONCAT`) return null when given a null argument. You can use the `NVL` function to return a value when a null occurs. For example, the expression `NVL(COMM,0)` returns 0 if `COMM` is null or the value of `COMM` if it is not null.

Most aggregate functions ignore nulls. For example, consider a query that averages the five values 1000, null, null, null, and 2000. Such a query ignores the nulls and calculates the average to be  $(1000+2000)/2 = 1500$ .

## Nulls with Comparison Operators

To test for nulls, use only the comparison operators `IS NULL` and `IS NOT NULL`. If you use any other operator with nulls and the result depends on the value of the null, the result is `UNKNOWN`. Because null represents a lack of data, a null cannot be equal or unequal to any value or to another null. However, Oracle considers two nulls to be equal when evaluating a `DECODE` expression. For syntax and additional information, see "[DECODE Expressions](#)" on page 5-12.

Oracle also considers two nulls to be equal if they appear in compound keys. That is, Oracle considers identical two compound keys containing nulls if all the non-null components of the keys are equal.

## Nulls in Conditions

A condition that evaluates to `UNKNOWN` acts almost like `FALSE`. For example, a `SELECT` statement with a condition in the `WHERE` clause that evaluates to `UNKNOWN` returns no rows. However, a condition evaluating to `UNKNOWN` differs from `FALSE` in that further operations on an `UNKNOWN` condition evaluation will evaluate to `UNKNOWN`. Thus, `NOT FALSE` evaluates to `TRUE`, but `NOT UNKNOWN` evaluates to `UNKNOWN`.

[Table 2-14](#) shows examples of various evaluations involving nulls in conditions. If the conditions evaluating to `UNKNOWN` were used in a `WHERE` clause of a `SELECT` statement, then no rows would be returned for that query.

**Table 2-14** *Conditions Containing Nulls*

If A is:	Condition	Evaluates to:
10	a IS NULL	FALSE
10	a IS NOT NULL	TRUE
NULL	a IS NULL	TRUE
NULL	a IS NOT NULL	FALSE
10	a = NULL	UNKNOWN
10	a != NULL	UNKNOWN
NULL	a = NULL	UNKNOWN
NULL	a != NULL	UNKNOWN
NULL	a = 10	UNKNOWN
NULL	a != 10	UNKNOWN

For the truth tables showing the results of logical expressions containing nulls, see [Table 3-6](#) on page 3-12, as well as [Table 3-7](#) and [Table 3-8](#).

## Pseudocolumns

A **pseudocolumn** behaves like a table column, but is not actually stored in the table. You can select from pseudocolumns, but you cannot insert, update, or delete their values. This section describes these pseudocolumns:

- [CURRVAL and NEXTVAL](#)
- [LEVEL](#)
- [ROWID](#)
- [ROWNUM](#)

### CURRVAL and NEXTVAL

A **sequence** is a schema object that can generate unique sequential values. These values are often used for primary and unique keys. You can refer to sequence values in SQL statements with these pseudocolumns:

`CURRVAL` returns the current value of a sequence.

`NEXTVAL` increments the sequence and returns the next value.

You must qualify `CURRVAL` and `NEXTVAL` with the name of the sequence:

```
sequence.CURRVAL
```

```
sequence.NEXTVAL
```

To refer to the current or next value of a sequence in the schema of another user, you must have been granted either `SELECT` object privilege on the sequence or `SELECT ANY SEQUENCE` system privilege, and you must qualify the sequence with the schema containing it:

```
schema.sequence.CURRVAL
```

```
schema.sequence.NEXTVAL
```

To refer to the value of a sequence on a remote database, you must qualify the sequence with a complete or partial name of a database link:

```
schema.sequence.CURRVAL@dblink
```

```
schema.sequence.NEXTVAL@dblink
```



**See Also:** ["Referring to Objects in Remote Databases"](#) on page 2-79 for more information on referring to database links.

### Where to Use Sequence Values

You can use `CURRVAL` and `NEXTVAL` in:

- The `SELECT` list of a `SELECT` statement that is not contained in a subquery, materialized view, or view
- The `SELECT` list of a subquery in an `INSERT` statement
- The `VALUES` clause of an `INSERT` statement
- The `SET` clause of an `UPDATE` statement

**Restrictions:** You cannot use `CURRVAL` and `NEXTVAL`:

- A subquery in a `DELETE`, `SELECT`, or `UPDATE` statement
- A query of a view or of a materialized view
- A `SELECT` statement with the `DISTINCT` operator
- A `SELECT` statement with a `GROUP BY` clause or `ORDER BY` clause
- A `SELECT` statement that is combined with another `SELECT` statement with the `UNION`, `INTERSECT`, or `MINUS` set operator
- The `WHERE` clause of a `SELECT` statement
- `DEFAULT` value of a column in a `CREATE TABLE` or `ALTER TABLE` statement
- The condition of a `CHECK` constraint

Also, within a single SQL statement that uses `CURVAL` or `NEXTVAL`, all referenced `LONG` columns, updated tables, and locked tables must be located on the same database.

### How to Use Sequence Values

When you create a sequence, you can define its initial value and the increment between its values. The first reference to `NEXTVAL` returns the sequence's initial value. Subsequent references to `NEXTVAL` increment the sequence value by the defined increment and return the new value. Any reference to `CURRVAL` always returns the sequence's current value, which is the value returned by the last reference to `NEXTVAL`. Note that before you use `CURRVAL` for a sequence in your session, you must first initialize the sequence with `NEXTVAL`.

Within a single SQL statement, Oracle will increment the sequence only once per row. If a statement contains more than one reference to `NEXTVAL` for a sequence, Oracle increments the sequence once and returns the same value for all occurrences of `NEXTVAL`. If a statement contains references to both `CURRVAL` and `NEXTVAL`, Oracle increments the sequence and returns the same value for both `CURRVAL` and `NEXTVAL` regardless of their order within the statement.

A sequence can be accessed by many users concurrently with no waiting or locking. For information on sequences, see ["CREATE SEQUENCE"](#) on page 4-371.

**Example 1** This example selects the current value of the employee sequence:

```
SELECT empseq.currval
       FROM DUAL;
```

**Example 2** This example increments the employee sequence and uses its value for a new employee inserted into the employee table:

```
INSERT INTO emp
       VALUES (empseq.nextval, 'LEWIS', 'CLERK',
              7902, SYSDATE, 1200, NULL, 20);
```

**Example 3** This example adds a new order with the next order number to the master order table. It then adds suborders with this number to the detail order table:

```
INSERT INTO master_order(orderno, customer, orderdate)
       VALUES (orderseq.nextval, 'Al's Auto Shop', SYSDATE);

INSERT INTO detail_order (orderno, part, quantity)
       VALUES (orderseq.currval, 'SPARKPLUG', 4);

INSERT INTO detail_order (orderno, part, quantity)
       VALUES (orderseq.currval, 'FUEL PUMP', 1);

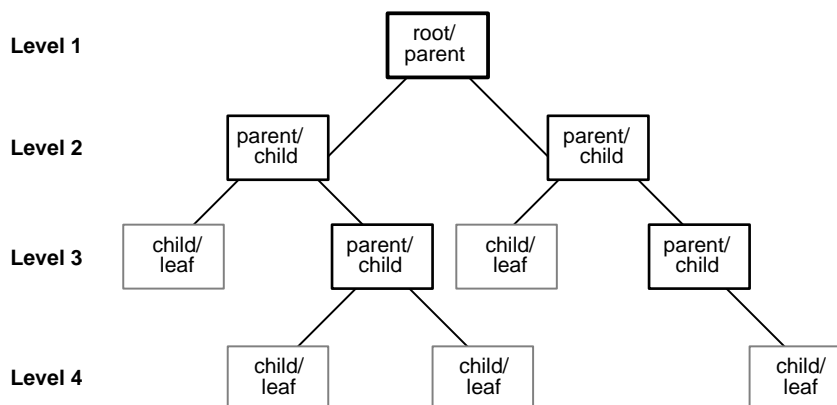
INSERT INTO detail_order (orderno, part, quantity)
       VALUES (orderseq.currval, 'TAILPIPE', 2);
```

## LEVEL

For each row returned by a hierarchical query, the `LEVEL` pseudocolumn returns 1 for a root node, 2 for a child of a root, and so on. A *root node* is the highest node within an inverted tree. A *child node* is any nonroot node. A *parent*

*node* is any node that has children. A *leaf node* is any node without children. Figure 2–2 shows the nodes of an inverted tree with their `LEVEL` values.

**Figure 2–2 Hierarchical Tree**



To define a hierarchical relationship in a query, you must use the `START WITH` and `CONNECT BY` clauses.

**See also:** ["SELECT and Subqueries"](#) on page 7-569 for more information on using the `LEVEL` pseudocolumn.

## ROWID

For each row in the database, the `ROWID` pseudocolumn returns a row's address. Oracle8i rowid values contain information necessary to locate a row:

- the data object number of the object
- which data block in the datafile
- which row in the data block (first row is 0)
- which datafile (first file is 1). The file number is relative to the tablespace.

Usually, a rowid value uniquely identifies a row in the database. However, rows in different tables that are stored together in the same cluster can have the same rowid.

Values of the `ROWID` pseudocolumn have the datatype `ROWID` or `UROWID`.

**See Also:** ["ROWID Datatype"](#) on page 2-26 and ["UROWID Datatype"](#) on page 2-27.

Rowid values have several important uses:

- They are the fastest way to access a single row.
- They can show you how a table's rows are stored.
- They are unique identifiers for rows in a table.

You should not use ROWID as a table's primary key. If you delete and reinsert a row with the Import and Export utilities, for example, its rowid may change. If you delete a row, Oracle may reassign its rowid to a new row inserted later.

Although you can use the ROWID pseudocolumn in the SELECT and WHERE clause of a query, these pseudocolumn values are not actually stored in the database. You cannot insert, update, or delete a value of the ROWID pseudocolumn.

**Example** This statement selects the address of all rows that contain data for employees in department 20:

```
SELECT ROWID, ename
       FROM emp
       WHERE deptno = 20;
```

```
ROWID          ENAME
-----
AAAAqYAABAAAEPvAAA SMITH
AAAAqYAABAAAEPvAAD JONES
AAAAqYAABAAAEPvAAH SCOTT
AAAAqYAABAAAEPvAAK ADAMS
AAAAqYAABAAAEPvAAM FORD
```

## ROWNUM

For each row returned by a query, the ROWNUM pseudocolumn returns a number indicating the order in which Oracle selects the row from a table or set of joined rows. The first row selected has a ROWNUM of 1, the second has 2, and so on.

You can use ROWNUM to limit the number of rows returned by a query, as in this example:

```
SELECT * FROM emp WHERE ROWNUM < 10;
```

If an `ORDER BY` clause follows `ROWNUM` in the same query, the rows will be reordered by the `ORDER BY` clause. The results can vary depending on the way the rows are accessed. For example, if the `ORDER BY` clause causes Oracle to use an index to access the data, Oracle may retrieve the rows in a different order than without the index. Therefore, the following statement will not have the same effect as the preceding example:

```
SELECT * FROM emp WHERE ROWNUM < 11 ORDER BY empno;
```

If you embed the `ORDER BY` clause in a subquery and place the `ROWNUM` condition in the top-level query, you can force the `ROWNUM` condition to be applied after the ordering of the rows. For example, the following query returns the 10 smallest employee numbers. This is sometimes referred to as a "top-N query":

```
SELECT * FROM
  (SELECT empno FROM emp ORDER BY empno)
 WHERE ROWNUM < 11;
```

In the preceding example, the `ROWNUM` values are those of the top-level `SELECT` statement, so they are generated after the rows have already been ordered by `EMPNO` in the subquery.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals* for more information about top-N queries.

Conditions testing for `ROWNUM` values greater than a positive integer are always false. For example, this query returns no rows:

```
SELECT * FROM emp
  WHERE ROWNUM > 1;
```

The first row fetched is assigned a `ROWNUM` of 1 and makes the condition false. The second row to be fetched is now the first row and is also assigned a `ROWNUM` of 1 and makes the condition false. All rows subsequently fail to satisfy the condition, so no rows are returned.

You can also use `ROWNUM` to assign unique values to each row of a table, as in this example:

```
UPDATE tabx
  SET col1 = ROWNUM;
```

---

---

**Note:** Using `ROWNUM` in a query can affect view optimization. For more information, see *Oracle8i Concepts*.

---

---

## Comments

You can associate comments with SQL statements and schema objects.

### Comments Within SQL Statements

Comments within SQL statements do not affect the statement execution, but they may make your application easier for you to read and maintain. You may want to include a comment in a statement that describes the statement's purpose within your application.

A comment can appear between any keywords, parameters, or punctuation marks in a statement. You can include a comment in a statement using either of these means:

- Begin the comment with a slash and an asterisk (`/*`). Proceed with the text of the comment. This text can span multiple lines. End the comment with an asterisk and a slash (`*/`). The opening and terminating characters need not be separated from the text by a space or a line break.
- Begin the comment with `--` (two hyphens). Proceed with the text of the comment. This text cannot extend to a new line. End the comment with a line break.

A SQL statement can contain multiple comments of both styles. The text of a comment can contain any printable characters in your database character set.

---

---

**Note:** You cannot use these styles of comments between SQL statements in a SQL script. Use the SQL\*Plus `REMARK` command for this purpose. For information on these statements, see *SQL\*Plus User's Guide and Reference*.

---

---

**Example** These statements contain many comments:

```
SELECT ename, sal + NVL(comm, 0), job, loc
/* Select all employees whose compensation is
greater than that of Jones.*/
FROM emp, dept
/*The DEPT table is used to get the department name.*/
```

```

WHERE emp.deptno = dept.deptno
      AND sal + NVL(comm,0) > /* Subquery:          */
      (SELECT sal + NLV(comm,0)
       FROM emp
       WHERE ename = 'JONES');
/* total compensation is sal + comm */

SELECT ename,           -- select the name
       sal + NVL(comm, 0), -- total compensation
       job,             -- job
       loc              -- and city containing the office
FROM emp,              -- of all employees
     dept
WHERE emp.deptno = dept.deptno
      AND sal + NVL(comm, 0) > -- whose compensation
                                   -- is greater than
      (SELECT sal + NVL(comm,0) -- the compensation
       FROM emp
       WHERE ename = 'JONES'); -- of Jones.

```

## Comments on Schema Objects

You can associate a comment with a table, view, materialized view, or column using the `COMMENT` command described in [Chapter 7, "SQL Statements"](#). Comments associated with schema objects are stored in the data dictionary.

## Hints

You can use comments in a SQL statement to pass instructions, or *hints*, to the Oracle optimizer. The optimizer uses these hints as suggestions for choosing an execution plan for the statement.

A statement block can have only one comment containing hints, and that comment must follow the `SELECT`, `UPDATE`, `INSERT`, or `DELETE` keyword. The syntax below shows hints contained in both styles of comments that Oracle supports within a statement block.

```
{DELETE|INSERT|SELECT|UPDATE} /*+ hint [text] [hint[text]]... */
```

or

```
{DELETE|INSERT|SELECT|UPDATE} --+ hint [text] [hint[text]]...
```

where

DELETE	is a DELETE, INSERT, SELECT, or UPDATE keyword that begins a statement block. Comments containing hints can appear only after these keywords.
INSERT	
SELECT	
UPDATE	
+	is a plus sign that causes Oracle to interpret the comment as a list of hints. The plus sign must follow immediately after the comment delimiter (no space is permitted).
<i>hint</i>	is one of the hints discussed in this section and in <i>Oracle8i Designing and Tuning for Performance</i> . The space between the plus sign and the hint is optional. If the comment contains multiple hints, separate the hints by at least one space.
<i>text</i>	is other commenting text that can be interspersed with the hints.

Table 2–15 lists hint syntax and descriptions.

**See Also:** *Oracle8i Designing and Tuning for Performance* and *Oracle8i Concepts* for more information on hints.

**Table 2–15 Hint Syntax and Descriptions**

Hint Syntax	Description
<b>Optimization Approaches and Goals</b>	
<code>/*+ ALL_ROWS */</code>	Explicitly chooses the cost-based approach to optimize a statement block with a goal of best throughput (that is, minimum total resource consumption).
<code>/*+ CHOOSE */</code>	Causes the optimizer to choose between the rule-based approach and the cost-based approach for a SQL statement based on the presence of statistics for the tables accessed by the statement.
<code>/*+ FIRST_ROWS */</code>	Explicitly chooses the cost-based approach to optimize a statement block with a goal of best response time (minimum resource usage to return first row).
<code>/*+ RULE */</code>	Explicitly chooses rule-based optimization for a statement block.



**Table 2–15 (Cont.) Hint Syntax and Descriptions**

Hint Syntax	Description
<b>Access Methods</b>	
<code>/*+ AND_EQUAL(table index) */</code>	Explicitly chooses an execution plan that uses an access path that merges the scans on several single-column indexes.
<code>/*+ CLUSTER(table) */</code>	Explicitly chooses a cluster scan to access the specified table.
<code>/*+ FULL(table) */</code>	Explicitly chooses a full table scan for the specified table.
<code>/*+ HASH(table) */</code>	Explicitly chooses a hash scan to access the specified table.
<code>/*+ HASH_AJ(table) */</code>	Transforms a <code>NOT IN</code> subquery into a hash anti-join to access the specified table.
<code>/*+ HASH_SJ(table) */</code>	Transforms a <code>NOT IN</code> subquery into a hash semi-join to access the specified table.
<code>/*+ INDEX(table index) */</code>	Explicitly chooses an index scan for the specified table.
<code>/*+ INDEX_ASC(table index) */</code>	Explicitly chooses an ascending-range index scan for the specified table.
<code>/*+ INDEX_COMBINE(table index) */</code>	If no indexes are given as arguments for the <code>INDEX_COMBINE</code> hint, the optimizer uses whatever Boolean combination of bitmap indexes has the best cost estimate. If particular indexes are given as arguments, the optimizer tries to use some Boolean combination of those particular bitmap indexes.
<code>/*+ INDEX_DESC(table index) */</code>	Explicitly chooses a descending-range index scan for the specified table.
<code>/*+ INDEX_FFS(table index) */</code>	Causes a fast full index scan to be performed rather than a full table scan.
<code>/*+ MERGE_AJ(table) */</code>	Transforms a <code>NOT IN</code> subquery into a merge anti-join to access the specified table.
<code>/*+ MERGE_SJ(table) */</code>	Transforms a correlated <code>EXISTS</code> subquery into a merge semi-join to access the specified table.
<code>/*+ NO_EXPAND */</code>	Prevents the optimizer from considering <code>OR</code> expansion for queries having <code>OR</code> or <code>IN</code> conditions in the <code>WHERE</code> clause.
<code>/*+ NO_INDEX(table index) */</code>	Instructs the optimizer not to consider a scan on the specified index or indexes. If no indexes are specified, the optimizer does not consider a scan on any index defined on the table.
<code>/*+ NOREWRITE */</code>	Disables query rewrite for the query block, overriding a <code>TRUE</code> setting of the <code>QUERY_REWRITE_ENABLED</code> parameter.

**Table 2–15 (Cont.) Hint Syntax and Descriptions**

Hint Syntax	Description
<code>/*+ NO_UNNEST */</code>	Prohibits the unnesting of the subquery. The hints <code>HASH_SJ</code> , <code>HASH_AJ</code> , <code>MERGE_SJ</code> , and <code>MERGE_AJ</code> take precedence over this hint.
<code>/*+ ORDERED_PREDICATES */</code>	Forces the optimizer to preserve the order of predicate evaluation (except predicates used in index keys), as specified in the <code>WHERE</code> clause of <code>SELECT</code> statements.
<code>/*+ REWRITE (view [,...]) */</code>	Enforces query rewrite. If you specify a view list and the list contains an eligible materialized view, Oracle will use that view regardless of the cost. No views outside of the list are considered. If you do not specify a view list, Oracle will search for an eligible materialized view and always use it regardless of the cost.
<code>/*+ ROWID(table) */</code>	Explicitly chooses a table scan by rowid for the specified table.
<code>/*+ UNNEST */</code>	Forces the unnesting of the subquery if possible.
<code>/*+ USE_CONCAT */</code>	Forces combined <code>OR</code> conditions in the <code>WHERE</code> clause of a query to be transformed into a compound query using the <code>UNION ALL</code> set operator.
<b>Join Orders</b>	
<code>/*+ ORDERED */</code>	Causes Oracle to join tables in the order in which they appear in the <code>FROM</code> clause. This hint overrides the <code>LEADING</code> hint.
<code>/*+ STAR */</code>	Forces the large table to be joined last using a nested-loops join on the index.
<b>Join Operations</b>	
<code>/*+ DRIVING_SITE(table) */</code>	Forces query execution to be done at a different site from that selected by Oracle.
<code>/*+ LEADING (table) */</code>	Causes Oracle to use the specified <i>table</i> as the first table in the join order. The <code>ORDERED</code> hint overrides this hint.
<code>/*+ USE_HASH(table) */</code>	Causes Oracle to join each specified table with another row source with a hash join.
<code>/*+ USE_MERGE(table) */</code>	Causes Oracle to join each specified table with another row source with a sort-merge join.
<code>/*+ USE_NL(table) */</code>	Causes Oracle to join each specified table to another row source with a nested-loops join using the specified table as the inner table.

**Table 2–15 (Cont.) Hint Syntax and Descriptions**

Hint Syntax	Description
<b>Parallel Execution</b>	
<b>Note:</b> Oracle ignores parallel hints on a temporary table. For more information on temporary tables, see "CREATE TABLE" on page 4-381 and <i>Oracle8i Concepts</i> .	
<code>/**+ APPEND */</code>	Specifies that data is simply appended (or not) to a table; existing free space is not used. Use these hints only following the <code>INSERT</code> keyword.
<code>/**+ NOAPPEND */</code>	
<code>/**+ NOPARALLEL(table) */</code>	Disables parallel scanning of a table, even if the table was created with a <code>PARALLEL</code> clause.  <b>Restriction:</b> You cannot parallelize a query involving a nested table.
<code>/**+ PARALLEL(table)</code> <code>/**+ PARALLEL(table, integer) */</code>	Lets you specify parallel execution of DML and queries on the table; <code>integer</code> specifies the desired degree of parallelism, which is the number of parallel threads that can be used for the operation. Each parallel thread may use one or two parallel execution servers. If you do not specify <code>integer</code> , Oracle computes a value using the <code>PARALLEL_THREADS_PER_CPU</code> parameter. If no parallel hint is specified, Oracle uses the existing degree of parallelism for the table.  DELETE, INSERT, and UPDATE operations are considered for parallelization only if the session is in a <code>PARALLEL DML</code> enabled mode. (Use <code>ALTER SESSION ENABLE PARALLEL DML</code> to enter this mode.)
<code>/**+ PARALLEL_INDEX</code>	Allows you to parallelize fast full index scans for partitioned and nonpartitioned indexes that have the <code>PARALLEL</code> attribute.
<code>/**+ PQ_DISTRIBUTE</code> <code>(table, outer_distribution,</code> <code>inner_distribution) */</code>	Specifies how rows of joined tables should be distributed between producer and consumer query servers. The four possible distribution methods are <code>NONE</code> , <code>HASH</code> , <code>BROADCAST</code> , and <code>PARTITION</code> . However, only a subset of the combinations of outer and inner distributions are valid. For the permitted combinations of distributions for the outer and inner join tables, see <i>Oracle8i Designing and Tuning for Performance</i> .
<code>/**+ NOPARALLEL_INDEX */</code>	Overrides a <code>PARALLEL</code> attribute setting on an index.
<b>Other Hints</b>	
<code>/**+ CACHE */</code>	Specifies that the blocks retrieved for the table in the hint are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed.

**Table 2–15 (Cont.) Hint Syntax and Descriptions**

<b>Hint Syntax</b>	<b>Description</b>
<code>/*+ NOCACHE */</code>	Specifies that the blocks retrieved for this table are placed at the least recently used end of the LRU list in the buffer cache when a full table scan is performed.
<code>/*+ MERGE(table) */</code>	Causes Oracle to evaluate complex views or subqueries before the surrounding query.
<code>/*+ NO_MERGE(table) */</code>	Causes Oracle not to merge mergeable views.
<code>/*+ PUSH_JOIN_PRED(table) */</code>	Causes the optimizer to evaluate, on a cost basis, whether to push individual join predicates into the view.
<code>/*+ NO_PUSH_JOIN_PRED(table) */</code>	Prevents pushing of a join predicate into the view.
<code>/*+ PUSH_SUBQ */</code>	Causes nonmerged subqueries to be evaluated at the earliest possible place in the execution plan.
<code>/*+ STAR_TRANSFORMATION */</code>	Makes the optimizer use the best plan in which the transformation has been used.

## Database Objects

Oracle recognizes objects that are associated with a particular schema and objects that are not associated with a particular schema, as described in the sections that follow.

### Schema Objects

A **schema** is a collection of logical structures of data, or schema objects. A schema is owned by a database user and has the same name as that user. Each user owns a single schema. Schema objects can be created and manipulated with SQL and include the following types of objects:

- clusters
- database links
- database triggers
- dimensions
- external procedure libraries
- index-organized tables
- indexes

- indextypes
- Java classes, Java resources, Java sources
- materialized views / snapshots
- materialized view logs / snapshot logs
- object tables
- object types
- object views
- operators
- packages
- sequences
- stored functions, stored procedures
- synonyms
- tables
- views

## Nonschema Objects

Other types of objects are also stored in the database and can be created and manipulated with SQL but are not contained in a schema:

- contexts
- directories
- profiles
- roles
- rollback segments
- tablespaces
- users

In this reference, each type of object is briefly defined in [Chapter 7, "SQL Statements"](#), in the section describing the statement that creates the database object. These statements begin with the keyword `CREATE`. For example, for the definition of a cluster, see "[CREATE CLUSTER](#)" on page 7-254. For an overview of database objects, see *Oracle8i Concepts*.

You must provide names for most types of schema objects when you create them. These names must follow the rules listed in the following sections.

## Parts of Schema Objects

Some schema objects are made up of parts that you can or must name, such as:

- columns in a table or view
- index and table partitions and subpartitions
- integrity constraints on a table
- packaged procedures, packaged stored functions, and other objects stored within a package

### Partitioned Tables and Indexes

Tables and indexes can be partitioned. When partitioned, these schema objects consist of a number of parts called *partitions*, all of which have the same logical attributes. For example, all partitions in a table share the same column and constraint definitions, and all partitions in an index share the same index columns.

When you partition a table or index using the range method, you specify a maximum value for the partitioning key column(s) for each partition. When you partition a table or index using the hash method, you instruct Oracle to distribute the rows of the table into partitions based on a system-defined hash function on the partitioning key column(s). When you partition a table or index using the composite-partitioning method, you specify ranges for the partitions, and Oracle distributes the rows in each partition into one or more hash subpartitions based on a hash function. Each subpartition of a table or index partitioned using the composite method has the same logical attributes.

### Partition-Extended and Subpartition-Extended Table Names

Partition-extended and subpartition-extended table names let you perform some partition-level and subpartition-level operations, such as deleting all rows from a partition or subpartition, on only one partition or subpartition. Without extended table names, such operations would require that you specify a predicate (`WHERE` clause). For range-partitioned tables, trying to phrase a partition-level operation with a predicate can be cumbersome, especially when the range partitioning key uses more than one column. For hash partitions and subpartitions, using a predicate is more difficult still, because these partitions and subpartitions are based on a system-defined hash function.

Partition-extended table names let you use partitions as if they were tables. An advantage of this method, which is most useful for range-partitioned tables, is that you can build partition-level access control mechanisms by granting (or revoking) privileges on these views to (or from) other users or roles. To use a partition as a table, create a view by selecting data from a single partition, and then use the view as a table.

You can specify partition-extended or subpartition-extended table names for the following DML statements:

- DELETE
- INSERT
- LOCK TABLE
- SELECT
- UPDATE

---



---


**Note:** For application portability and ANSI syntax compliance, Oracle strongly recommends that you use views to insulate applications from this Oracle proprietary extension.

---



---

**Syntax** The basic syntax for using partition-extended and subpartition-extended table names is:



**Restrictions** Currently, the use of partition-extended and subpartition-extended table names has the following restrictions:

- No remote tables: A partition-extended or subpartition-extended table name cannot contain a database link (dblink) or a synonym that translates to a table with a dblink. To use remote partitions and subpartitions, create a view at the remote site that uses the extended table name syntax and then refer to the remote view.

- No synonyms: A partition or subpartition extension must be specified with a base table. You cannot use synonyms, views, or any other objects.

**Example** In the following statement, SALES is a partitioned table with partition JAN97. You can create a view of the single partition JAN97, and then use it as if it were a table. This example deletes rows from the partition.

```
CREATE VIEW sales_jan97 AS
    SELECT * FROM sales PARTITION (jan97);
DELETE FROM sales_jan97 WHERE amount < 0;
```

## Schema Object Names and Qualifiers

This section provides:

- rules for naming schema objects and schema object location qualifiers
- guidelines for naming schema objects and qualifiers

## Schema Object Naming Rules

The following rules apply when naming schema objects:

1. Names must be from 1 to 30 bytes long with these exceptions:
  - Names of databases are limited to 8 bytes.
  - Names of database links can be as long as 128 bytes.
2. Names cannot contain quotation marks.
3. Names are not case sensitive.
4. A name must begin with an alphabetic character from your database character set unless surrounded by double quotation marks.
5. Names can contain only alphanumeric characters from your database character set and the underscore (\_), dollar sign (\$), and pound sign (#). Oracle strongly discourages you from using \$ and #. Names of database links can also contain periods (.) and "at" signs (@).

If your database character set contains multibyte characters, Oracle recommends that each name for a user or a role contain at least one single-byte character.



---

---

**Note:** You cannot use special characters from European or Asian character sets in a database name, global database name, or database link names. For example, characters with an umlaut are not allowed.

---

---

6. A name cannot be an Oracle reserved word. [Appendix C, "Oracle Reserved Words"](#), lists all Oracle reserved words.

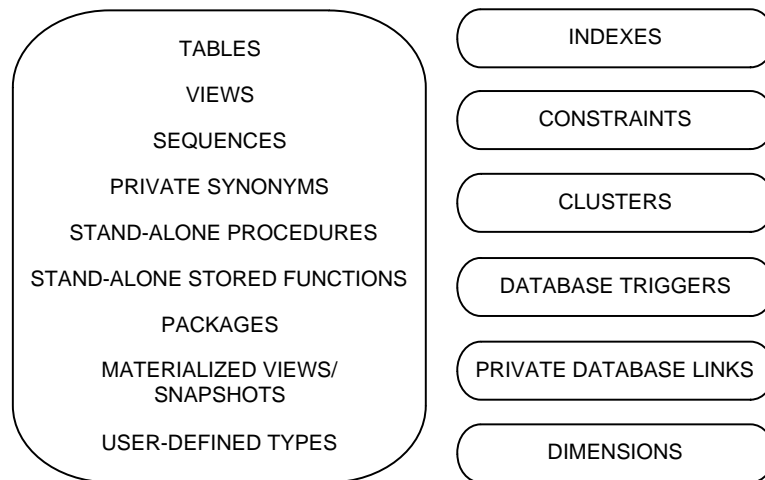
Depending on the Oracle product you plan to use to access a database object, names might be further restricted by other product-specific reserved words. For a list of a product's reserved words, see the manual for the specific product, such as *PL/SQL User's Guide and Reference*.

7. Do not use the word DUAL as a name for an object or part. DUAL is the name of a dummy table.
8. The Oracle SQL language contains other words that have special meanings. These words include datatypes (see "[Datatypes](#)" on page 2-9), function names (see "[SQL Functions](#)" on page 4-2), and keywords (the uppercase words in SQL statements, such as DIMENSION, SEGMENT, ALLOCATE, DISABLE, and so forth). These words are not reserved. However, Oracle uses them internally. Therefore, if you use these words as names for objects and object parts, your SQL statements may be more difficult to read and may lead to unpredictable results.

In particular, do not use words beginning with "SYS\_" as schema object names, and do not use the names of SQL built-in functions for the names of schema objects or user-defined functions.


9. Within a namespace, no two objects can have the same name.

The following figure shows the namespaces for schema objects. Each box is a namespace. Tables and views are in the same namespace. Therefore, a table and a view in the same schema cannot have the same name. However, tables and indexes are in different namespaces. Therefore, a table and an index in the same schema can have the same name.



Each schema in the database has its own namespaces for the objects it contains. This means, for example, that two tables in different schemas are in different namespaces and can have the same name.

The following figure shows the namespaces for nonschema objects. Because the objects in these namespaces are not contained in schemas, these namespaces span the entire database.



10. Columns in the same table or view cannot have the same name. However, columns in different tables or views can have the same name.
11. Procedures or functions contained in the same package can have the same name, provided that their arguments are not of the same number and datatypes. Creating multiple procedures or functions with the same name in the same package with different arguments is called *overloading* the procedure or function.

12. A name can be enclosed in double quotation marks. Such names can contain any combination of characters, including spaces, ignoring rules 3 through 7 in this list. This exception is allowed for portability, but Oracle recommends that you do not break rules 3 through 7.

If you give a schema object a name enclosed in double quotation marks, you must use double quotation marks whenever you refer to the object.

Enclosing a name in double quotes allows it to:

- Contain spaces
- Be case sensitive
- Begin with a character other than an alphabetic character, such as a numeric character
- Contain characters other than alphanumeric characters and `_`, `$`, and `#`
- Be a reserved word

By enclosing names in double quotation marks, you can give the following names to different objects in the same namespace:

```
emp
"emp"
"Emp"
"EMP "
```

Note that Oracle interprets the following names the same, so they cannot be used for different objects in the same namespace:

```
emp
EMP
"EMP "
```

If you give a user or password a quoted name, the name cannot contain lowercase letters.

Database link names cannot be quoted.

## Schema Object Naming Examples

The following examples are valid schema object names:

```
ename
horse
scott.hiredate
```

```
"EVEN THIS & THAT!"  
a_very_long_and_valid_name
```

Although column aliases, table aliases, usernames, and passwords are not objects or parts of objects, they must also follow these naming rules with these exceptions:

- Column aliases and table aliases exist only for the execution of a single SQL statement and are not stored in the database, so rule 12 does not apply to them.
- Passwords do not have namespaces, so rule 9 does not apply to them.
- Do not use quotation marks to make usernames and passwords case sensitive. For additional rules for naming users and passwords, see "CREATE USER" on page 7-451.

## Schema Object Naming Guidelines

Here are several helpful guidelines for naming objects and their parts:

- Use full, descriptive, pronounceable names (or well-known abbreviations).
- Use consistent naming rules.
- Use the same name to describe the same entity or attribute across tables.

When naming objects, balance the objective of keeping names short and easy to use with the objective of making names as descriptive as possible. When in doubt, choose the more descriptive name, because the objects in the database may be used by many people over a period of time. Your counterpart ten years from now may have difficulty understanding a database with a name like PMDD instead of PAYMENT\_DUE\_DATE.

Using consistent naming rules helps users understand the part that each table plays in your application. One such rule might be to begin the names of all tables belonging to the FINANCE application with FIN\_.

Use the same names to describe the same things across tables. For example, the department number columns of the sample EMP and DEPT tables are both named DEPTNO.


## Referring to Schema Objects and Parts

This section tells you how to refer to schema objects and their parts in the context of a SQL statement. This section shows you:

- the general syntax for referring to an object

- how Oracle resolves a reference to an object
- how to refer to objects in schemas other than your own
- how to refer to objects in remote databases

The following diagram shows the general syntax for referring to an object or a part:



where:

*object* is the name of the object.

*schema* is the schema containing the object. The schema qualifier allows you to refer to an object in a schema other than your own. You must be granted privileges to refer to objects in other schemas. If you omit *schema*, Oracle assumes that you are referring to an object in your own schema.

Only schema objects can be qualified with *schema*. Schema objects are shown with list item 9 on page 2-73. Nonschema objects, also shown with list item 9 on page 2-73, cannot be qualified with *schema* because they are not schema objects. (An exception is public synonyms, which can optionally be qualified with "PUBLIC". The quotation marks are required.)

*part* is a part of the object. This identifier allows you to refer to a part of a schema object, such as a column or a partition of a table. Not all types of objects have parts.

*dblink* applies only when you are using Oracle's distributed functionality. This is the name of the database containing the object. The *dblink* qualifier lets you refer to an object in a database other than your local database. If you omit *dblink*, Oracle assumes that you are referring to an object in your local database. Not all SQL statements allow you to access objects on remote databases.

You can include spaces around the periods separating the components of the reference to the object, but it is conventional to omit them.

## How Oracle Resolves Schema Object References

When you refer to an object in a SQL statement, Oracle considers the context of the SQL statement and locates the object in the appropriate namespace. After locating the object, Oracle performs the statement's operation on the object. If the named object cannot be found in the appropriate namespace, Oracle returns an error.

The following example illustrates how Oracle resolves references to objects within SQL statements. Consider this statement that adds a row of data to a table identified by the name DEPT:

```
INSERT INTO dept
VALUES (50, 'SUPPORT', 'PARIS');
```

Based on the context of the statement, Oracle determines that DEPT can be:

- A table in your own schema
- A view in your own schema
- A private synonym for a table or view
- A public synonym

Oracle always attempts to resolve an object reference within the namespaces in your own schema before considering namespaces outside your schema. In this example, Oracle attempts to resolve the name DEPT as follows:

1. First, Oracle attempts to locate the object in the namespace in your own schema containing tables, views, and private synonyms. If the object is a private synonym, Oracle locates the object for which the synonym stands. This object could be in your own schema, another schema, or on another database. The object could also be another synonym, in which case Oracle locates the object for which this synonym stands.
2. If the object is in the namespace, Oracle attempts to perform the statement on the object. In this example, Oracle attempts to add the row of data to DEPT. If the object is not of the correct type for the statement, Oracle returns an error. In this example, DEPT must be a table, view, or a private synonym resolving to a table or view. If DEPT is a sequence, Oracle returns an error.
3. If the object is not in any namespace searched in thus far, Oracle searches the namespace containing public synonyms. If the object is in that namespace, Oracle attempts to perform the statement on it. If the object is not of the correct type for the statement, Oracle returns an error. In this example, if DEPT is a public synonym for a sequence, Oracle returns an error.

## Referring to Objects in Other Schemas

To refer to objects in schemas other than your own, prefix the object name with the schema name:

```
schema.object
```

For example, this statement drops the EMP table in the schema SCOTT:

```
DROP TABLE scott.emp
```

## Referring to Objects in Remote Databases

To refer to objects in databases other than your local database, follow the object name with the name of the database link to that database. A database link is a schema object that causes Oracle to connect to a remote database to access an object there. This section tells you:

- How to create database links
- How to use database links in your SQL statements

### Creating Database Links

You create a database link with the `CREATE DATABASE LINK` statement described in [Chapter 7, "SQL Statements"](#). The statement allows you to specify this information about the database link:


- The name of the database link
- The database connect string to access the remote database
- The username and password to connect to the remote database

Oracle stores this information in the data dictionary.

**Database Link Names** When you create a database link, you must specify its name. Database link names are different from names of other types of objects. They can be as long as 128 bytes and can contain periods (.) and the "at" sign (@).

The name that you give to a database link must correspond to the name of the database to which the database link refers and the location of that database in the hierarchy of database names. The following syntax diagram shows the form of the name of a database link:

**dblink::=**



**where:**

- database** should specify *name* portion of the *global name* of the remote database to which the database link connects. This global name is stored in the data dictionary of the remote database; you can see this name in the GLOBAL\_NAME view.
- domain** should specify the *domain* portion of the global name of the remote database to which the database link connects. If you omit *domain* from the name of a database link, Oracle qualifies the database link name with the domain of your local database as it currently exists in the data dictionary.
- connect\_descriptor** allows you to further qualify a database link. Using connect descriptors, you can create multiple database links to the same database. For example, you can use connect descriptors to create multiple database links to different instances of the Oracle Parallel Server that access the same database.

The combination *database.domain* is sometimes called the "service name".

**See Also:** *Net8 Administrator's Guide*.

**Username and Password** Oracle uses the username and password to connect to the remote database. The username and password for a database link are optional.

**Database Connect String** The database connect string is the specification used by Net8 to access the remote database. For information on writing database connect strings, see the Net8 documentation for your specific network protocol. The database string for a database link is optional.



## Referring to Database Links

Database links are available only if you are using Oracle's distributed functionality. When you issue a SQL statement that contains a database link, you can specify the database link name in one of these forms:

- |                 |   |
|-----------------|---|
| <i>complete</i> | is the complete database link name as stored in the data dictionary, including the <i>database</i> , <i>domain</i> , and optional <i>connect_descriptor</i> components. |
| <i>partial</i>  | is the <i>database</i> and optional <i>connect_descriptor</i> components, but not the <i>domain</i> component.  |

Oracle performs these tasks before connecting to the remote database:

1. If the database link name specified in the statement is partial, Oracle expands the name to contain the domain of the local database as found in the global database name stored in the data dictionary. (You can see the current global database name in the GLOBAL\_NAME data dictionary view.)
2. Oracle first searches for a private database link in your own schema with the same name as the database link in the statement. Then, if necessary, it searches for a public database link with the same name.
  - Oracle always determines the username and password from the first matching database link (either private or public). If the first matching database link has an associated username and password, Oracle uses it. If it does not have an associated username and password, Oracle uses your current username and password.
  - If the first matching database link has an associated database string, Oracle uses it. If not, Oracle searches for the next matching (public) database link. If no matching database link is found, or if no matching link has an associated database string, Oracle returns an error.
3. Oracle uses the database string to access the remote database. After accessing the remote database, if the value of the GLOBAL\_NAMES parameter is TRUE, Oracle verifies that the *database.domain* portion of the database link name matches the complete global name of the remote database. If this condition is true, Oracle proceeds with the connection, using the username and password chosen in Step 2. If not, Oracle returns an error.
4. If the connection using the database string, username, and password is successful, Oracle attempts to access the specified object on the remote database using the rules for resolving object references and referring to objects in other schemas discussed earlier in this section.

You can disable the requirement that the *database.domain* portion of the database link name must match the complete global name of the remote database by setting to FALSE the initialization parameter GLOBAL\_NAMES or the GLOBAL\_NAMES parameter of the ALTER SYSTEM or ALTER SESSION statement.

**See Also:** *Oracle8i Distributed Database Systems* for more information on remote name resolution.

## Referencing Object Type Attributes and Methods

To reference object type attributes or methods in a SQL statement, you must fully qualify the reference with a table alias. Consider the following example:

```
CREATE TYPE person AS OBJECT
  (ssno VARCHAR(20),
   name VARCHAR (10));

CREATE TABLE emptab (pinfo person);
```

In a SQL statement, reference to the SSNO attribute must be fully qualified using a table alias, as illustrated below:

```
SELECT e.pinfo.ssno FROM emptab e;

UPDATE emptab e SET e.pinfo.ssno = '510129980'
  WHERE e.pinfo.name = 'Mike';
```

To reference an object type's member method that does not accept arguments, you must provide "empty" parentheses. For example, assume that AGE is a method in the person type that does not take arguments. In order to call this method in a SQL statement, you must provide empty parentheses as shows in this example:

```
SELECT e.pinfo.age() FROM emptab e
  WHERE e.pinfo.name = 'Mike';
```

**See Also:** *Oracle8i Concepts* for more information on user-defined datatypes.

---

# Operators

*With affection beaming in one eye, and calculation shining out of the other.*

Charles Dickens, *Martin Chuzzlewit*

An operator manipulates individual data items and returns a result. The data items are called *operands* or *arguments*. Operators are represented by special characters or by keywords. For example, the multiplication operator is represented by an asterisk (\*) and the operator that tests for nulls is represented by the keywords `IS NULL`.

This chapter discusses the following topics:

- [Unary and Binary Operators](#)
- [Precedence](#)
- [Arithmetic Operators](#)
- [Concatenation Operator](#)
- [Comparison Operators](#)
- [Logical Operators](#)
- [Set Operators](#)
- [Other Built-In Operators](#)
- [User-Defined Operators](#)

## Unary and Binary Operators

The two general classes of operators are:

unary	A unary operator operates on only one operand. A unary operator typically appears with its operand in this format: <code>operator operand</code>
binary	A binary operator operates on two operands. A binary operator appears with its operands in this format: <code>operand1 operator operand2</code>

Other operators with special formats accept more than two operands. If an operator is given a null operand, the result is always null. The only operator that does not follow this rule is concatenation ( `||` ).

## Precedence

*Precedence* is the order in which Oracle evaluates different operators in the same expression. When evaluating an expression containing multiple operators, Oracle evaluates operators with higher precedence before evaluating those with lower precedence. Oracle evaluates operators with equal precedence from left to right within an expression.

[Table 3–1](#) lists the levels of precedence among SQL operators from high to low. Operators listed on the same line have the same precedence.

**Table 3–1** SQL Operator Precedence

Operator	Operation
+, -	identity, negation
*, /	multiplication, division
+, -,	addition, subtraction, concatenation
=, !=, <, >, <=, >=, IS NULL, LIKE, BETWEEN, IN	comparison
NOT	exponentiation, logical negation
AND	conjunction
OR	disjunction

**Example** In the following expression, multiplication has a higher precedence than addition, so Oracle first multiplies 2 by 3 and then adds the result to 1.

`1+2*3`

You can use parentheses in an expression to override operator precedence. Oracle evaluates expressions inside parentheses before evaluating those outside.

SQL also supports set operators (UNION, UNION ALL, INTERSECT, and MINUS), which combine sets of rows returned by queries, rather than individual data items. All set operators have equal precedence.

## Arithmetic Operators

You can use an arithmetic operator in an expression to negate, add, subtract, multiply, and divide numeric values. The result of the operation is also a numeric value. Some of these operators are also used in date arithmetic. [Table 3-2](#) lists arithmetic operators.

**Table 3-2** *Arithmetic Operators*

Operator	Purpose	Example
+ -	Denotes a positive or negative expression. These are unary operators.	<pre>SELECT * FROM orders WHERE qty sold = -1; SELECT * FROM emp WHERE -sal &lt; 0;</pre>
* /	Multiplies, divides. These are binary operators.	<pre>UPDATE emp SET sal = sal * 1.1;</pre>
+ -	Adds, subtracts. These are binary operators.	<pre>SELECT sal + comm FROM emp WHERE SYSDATE - hiredate &gt; 365;</pre>

Do not use two consecutive minus signs (--) in arithmetic expressions to indicate double negation or the subtraction of a negative value. The characters -- are used to begin comments within SQL statements. You should separate consecutive minus signs with a space or a parenthesis.

**See Also:** ["Comments"](#) on page 2-62 for more information on comments within SQL statements.

## Concatenation Operator

The concatenation operator manipulates character strings. [Table 3-3](#) describes the concatenation operator.

**Table 3–3 Concatenation Operator**

Operator	Purpose	Example
	Concatenates character strings.	SELECT 'Name is '    ename FROM emp;

The result of concatenating two character strings is another character string. If both character strings are of datatype `CHAR`, the result has datatype `CHAR` and is limited to 2000 characters. If either string is of datatype `VARCHAR2`, the result has datatype `VARCHAR2` and is limited to 4000 characters. Trailing blanks in character strings are preserved by concatenation, regardless of the strings' datatypes.

On most platforms, the concatenation operator is two solid vertical bars, as shown in [Table 3–3](#). However, some IBM platforms use broken vertical bars for this operator. When moving SQL script files between systems having different character sets, such as between ASCII and EBCDIC, vertical bars might not be translated into the vertical bar required by the target Oracle environment. Oracle provides the `CONCAT` character function as an alternative to the vertical bar operator for cases when it is difficult or impossible to control translation performed by operating system or network utilities. Use this function in applications that will be moved between environments with differing character sets.

Although Oracle treats zero-length character strings as nulls, concatenating a zero-length character string with another operand always results in the other operand, so null can result only from the concatenation of two null strings. However, this may not continue to be true in future versions of Oracle. To concatenate an expression that might be null, use the `NVL` function to explicitly convert the expression to a zero-length string.

**See Also:** ["Character Datatypes"](#) on page 2-14 for more information on the differences between the `CHAR` and `VARCHAR2` datatypes.

**Example** This example creates a table with both `CHAR` and `VARCHAR2` columns, inserts values both with and without trailing blanks, and then selects these values and concatenates them. Note that for both `CHAR` and `VARCHAR2` columns, the trailing blanks are preserved.

```
CREATE TABLE tab1 (col1 VARCHAR2(6), col2 CHAR(6),
                  col3 VARCHAR2(6), col4 CHAR(6) );
```

Table created.

```
INSERT INTO tabl (col1, col2, col3, col4)
VALUES ('abc', 'def ', 'ghi ', 'jkl');
```

1 row created.

```
SELECT col1||col2||col3||col4 "Concatenation"
FROM tabl;
```

Concatenation

```
-----
abcdef  ghi   jkl
```

## Comparison Operators

Comparison operators compare one expression with another. The result of such a comparison can be TRUE, FALSE, or UNKNOWN. For information on conditions, see ["Conditions"](#) on page 5-14. [Table 3-4](#) lists comparison operators.

**Table 3-4 Comparison Operators**

Operator	Purpose	Example
=	Equality test.	SELECT * FROM emp WHERE sal = 1500;
!= ^= < > -!=	Inequality test. Some forms of the inequality operator may be unavailable on some platforms.	SELECT * FROM emp WHERE sal != 1500;
>	"Greater than" and "less than" tests.	SELECT * FROM emp WHERE sal > 1500;
<		SELECT * FROM emp WHERE sal < 1500;
>=	"Greater than or equal to" and "less than or equal to" tests.	SELECT * FROM emp WHERE sal >= 1500;
<=		SELECT * FROM emp WHERE sal <= 1500;

**Table 3–4 (Cont.) Comparison Operators**

<b>Operator</b>	<b>Purpose</b>	<b>Example</b>
IN	"Equal to any member of" test. Equivalent to "= ANY".	<pre>SELECT * FROM emp WHERE job IN ('CLERK', 'ANALYST'); SELECT * FROM emp WHERE sal IN (SELECT sal FROM emp WHERE deptno = 30);</pre>
NOT IN	Equivalent to "!=ALL". Evaluates to FALSE if any member of the set is NULL.	<pre>SELECT * FROM emp WHERE sal NOT IN (SELECT sal FROM emp WHERE deptno = 30); SELECT * FROM emp WHERE job NOT IN ('CLERK', 'ANALYST');</pre>
ANY SOME	Compares a value to each value in a list or returned by a query. Must be preceded by =, !=, >, <, <=, >=. Evaluates to FALSE if the query returns no rows.	<pre>SELECT * FROM emp WHERE sal = ANY (SELECT sal FROM emp WHERE deptno = 30);</pre>
ALL	Compares a value to every value in a list or returned by a query. Must be preceded by =, !=, >, <, <=, >=. Evaluates to TRUE if the query returns no rows.	<pre>SELECT * FROM emp WHERE sal &gt;= ALL ( 1400, 3000);</pre>
[NOT] BETWEEN x AND y	[Not] greater than or equal to x and less than or equal to y.	<pre>SELECT * FROM emp WHERE sal BETWEEN 2000 AND 3000;</pre>
EXISTS	TRUE if a subquery returns at least one row.	<pre>SELECT ename, deptno FROM dept WHERE EXISTS (SELECT * FROM emp WHERE dept.deptno = emp.deptno);</pre>



**Table 3–4 (Cont.) Comparison Operators**

Operator	Purpose	Example
x [NOT] LIKE y [ESCAPE 'z']	TRUE if x does [not] match the pattern y. Within y, the character "%" matches any string of zero or more characters except null. The character "_" matches any single character. Any character, excepting percent (%) and underbar ( ) may follow ESCAPE. A wildcard character is treated as a literal if preceded by the character designated as the escape character.	See " <a href="#">LIKE Operator</a> " on page 3-8. <pre>SELECT * FROM tabl WHERE coll LIKE 'A_C/%E%' ESCAPE '/';</pre>
IS [NOT] NULL	Tests for nulls. This is the only operator that you should use to test for nulls. See " <a href="#">Nulls</a> " on page 2-54.	<pre>SELECT ename, deptno FROM emp WHERE comm IS NULL;</pre>

Additional information on the NOT IN and LIKE operators appears in the sections that follow.

## NOT IN Operator

If any item in the list following a NOT IN operation is null, all rows evaluate to UNKNOWN (and no rows are returned). For example, the following statement returns the string 'TRUE' for each row:

```
SELECT 'TRUE'
FROM emp
WHERE deptno NOT IN (5,15);
```

However, the following statement returns no rows:

```
SELECT 'TRUE'
FROM emp
WHERE deptno NOT IN (5,15,null);
```


The above example returns no rows because the WHERE clause condition evaluates to:

```
deptno != 5 AND deptno != 15 AND deptno != null
```

Because all conditions that compare a null result in a null, the entire expression results in a null. This behavior can easily be overlooked, especially when the NOT IN operator references a subquery.

## LIKE Operator

The `LIKE` operator is used in character string comparisons with pattern matching. The syntax for a condition using the `LIKE` operator is shown in this diagram:



where:

- char1* is a value to be compared with a pattern. This value can have datatype `CHAR` or `VARCHAR2`.
- `NOT` logically inverts the result of the condition, returning `FALSE` if the condition evaluates to `TRUE` and `TRUE` if it evaluates to `FALSE`.
- char2* is the pattern to which *char1* is compared. The pattern is a value of datatype `CHAR` or `VARCHAR2` and can contain the special pattern matching characters `%` and `_`.
- `ESCAPE` identifies a single character as the escape character. The escape character can be used to cause Oracle to interpret `%` or `_` literally, rather than as a special character.  
  
If you wish to search for strings containing an escape character, you must specify this character twice. For example, if the escape character is `'/'`, to search for the string `'client/server'`, you must specify, `'client//server'`.

Whereas the equal (`=`) operator exactly matches one character value to another, the `LIKE` operator matches a portion of one character value to another by searching the first value for the pattern specified by the second. Note that blank padding is **not** used for `LIKE` comparisons.

With the `LIKE` operator, you can compare a value to a pattern rather than to a constant. The pattern must appear after the `LIKE` keyword. For example, you can issue the following query to find the salaries of all employees with names beginning with `'SM'`:

```
SELECT sal
   FROM emp
  WHERE ename LIKE 'SM%';
```

The following query uses the = operator, rather than the LIKE operator, to find the salaries of all employees with the name 'SM%':

```
SELECT sal
   FROM emp
  WHERE ename = 'SM%';
```

The following query finds the salaries of all employees with the name 'SM%'. Oracle interprets 'SM%' as a text literal, rather than as a pattern, because it *precedes* the LIKE operator:

```
SELECT sal
   FROM emp
  WHERE 'SM%' LIKE ename;
```

Patterns typically use special characters that Oracle matches with different characters in the value:

- An underscore ( `_` ) in the pattern matches exactly one character (as opposed to one byte in a multibyte character set) in the value.
- A percent sign ( `%` ) in the pattern can match zero or more characters (as opposed to bytes in a multibyte character set) in the value. Note that the pattern `'%'` cannot match a null.

**Case Sensitivity and Pattern Matching** Case is significant in all conditions comparing character expressions including the LIKE and equality (=) operators. You can use the UPPER() function to perform a case-insensitive match, as in this condition:

```
UPPER(ename) LIKE 'SM%'
```

**Pattern Matching on Indexed Columns** When LIKE is used to search an indexed column for a pattern, Oracle can use the index to improve the statement's performance if the leading character in the pattern is not "%" or "\_". In this case, Oracle can scan the index by this leading character. If the first character in the pattern is "%" or "\_", the index cannot improve the query's performance because Oracle cannot scan the index.

**Example 1** This condition is true for all ENAME values beginning with "MA":

```
ename LIKE 'MA%'
```

All of these ENAME values make the condition TRUE:

```
MARTIN, MA, MARK, MARY
```

Case is significant, so ENAME values beginning with "Ma," "ma," and "mA" make the condition FALSE.

**Example 2** Consider this condition:

```
ename LIKE 'SMITH_'
```

This condition is true for these ENAME values:

```
SMITHE, SMITHY, SMITHS
```

This condition is false for 'SMITH', since the special character "\_" must match exactly one character of the ENAME value.

**ESCAPE Option** You can include the actual characters "%" or "\_" in the pattern by using the ESCAPE option. The ESCAPE option identifies the escape character. If the escape character appears in the pattern before the character "%" or "\_" then Oracle interprets this character literally in the pattern, rather than as a special pattern matching character.

**Example:** To search for employees with the pattern 'A\_B' in their name:

```
SELECT ename
FROM emp
WHERE ename LIKE '%A\_B%' ESCAPE '\';
```

The ESCAPE option identifies the backslash (\) as the escape character. In the pattern, the escape character precedes the underscore (\_). This causes Oracle to interpret the underscore literally, rather than as a special pattern matching character.

**Patterns Without %** If a pattern does not contain the "%" character, the condition can be TRUE only if both operands have the same length.

**Example:** Consider the definition of this table and the values inserted into it:

```
CREATE TABLE freds (f CHAR(6), v VARCHAR2(6));
INSERT INTO freds VALUES ('FRED', 'FRED');
```

Because Oracle blank-pads CHAR values, the value of F is blank-padded to 6 bytes. V is not blank-padded and has length 4.

## Logical Operators

A logical operator combines the results of two component conditions to produce a single result based on them or to invert the result of a single condition. [Table 3–5](#) lists logical operators.

**Table 3–5 Logical Operators**

Operator	Function	Example
NOT	Returns TRUE if the following condition is FALSE. Returns FALSE if it is TRUE. If it is UNKNOWN, it remains UNKNOWN.	<pre>SELECT *   FROM emp  WHERE NOT (job IS NULL); SELECT *   FROM emp  WHERE NOT  (sal BETWEEN 1000 AND 2000);</pre>
AND	Returns TRUE if both component conditions are TRUE. Returns FALSE if either is FALSE. Otherwise returns UNKNOWN.	<pre>SELECT *   FROM emp  WHERE job = 'CLERK'  AND deptno = 10;</pre>
OR	Returns TRUE if either component condition is TRUE. Returns FALSE if both are FALSE. Otherwise returns UNKNOWN.	<pre>SELECT *   FROM emp  WHERE job = 'CLERK'  OR deptno = 10;</pre>

For example, in the WHERE clause of the following SELECT statement, the AND logical operator is used to ensure that only those hired before 1984 and earning more than \$1000 a month are returned:

```
SELECT *
  FROM emp
 WHERE hiredate < TO_DATE('01-JAN-1984', 'DD-MON-YYYY')
 AND sal > 1000;
```

## NOT Operator

[Table 3–6](#) shows the result of applying the NOT operator to a condition.

**Table 3–6 NOT Truth Table**

	TRUE	FALSE	UNKNOWN
NOT	FALSE	TRUE	UNKNOWN

## AND Operator

[Table 3–7](#) shows the results of combining two expressions with AND.

**Table 3–7 AND Truth Table**

AND	TRUE	FALSE	UNKNOWN
TRUE	TRUE	FALSE	UNKNOWN
FALSE	FALSE	FALSE	FALSE
UNKNOWN	UNKNOWN	FALSE	UNKNOWN

## OR Operator

[Table 3–8](#) shows the results of combining two expressions with OR.

**Table 3–8 OR Truth Table**

OR	TRUE	FALSE	UNKNOWN
TRUE	TRUE	TRUE	TRUE
FALSE	TRUE	FALSE	UNKNOWN
UNKNOWN	TRUE	UNKNOWN	UNKNOWN

## Set Operators

Set operators combine the results of two component queries into a single result. Queries containing set operators are called compound queries. [Table 3–9](#) lists SQL set operators.

**Table 3–9 Set Operators**

Operator	Returns
UNION	All rows selected by either query.
UNION ALL	All rows selected by either query, including all duplicates.
INTERSECT	All distinct rows selected by both queries.
MINUS	All distinct rows selected by the first query but not the second.

All set operators have equal precedence. If a SQL statement contains multiple set operators, Oracle evaluates them from the left to right if no parentheses explicitly specify another order.

The corresponding expressions in the select lists of the component queries of a compound query must match in number and datatype. If component queries select character data, the datatype of the return values are determined as follows:

- If both queries select values of datatype CHAR, the returned values have datatype CHAR.
- If either or both of the queries select values of datatype VARCHAR2, the returned values have datatype VARCHAR2.

**Examples** Consider these two queries and their results:

```
SELECT part
       FROM orders_list1;
```

```
PART
-----
SPARKPLUG
FUEL PUMP
FUEL PUMP
TAILPIPE
```

```
SELECT part
       FROM orders_list2;
```

```
PART
-----
CRANKSHAFT
TAILPIPE
TAILPIPE
```

The following examples combine the two query results with each of the set operators.

**UNION Example** The following statement combines the results with the UNION operator, which eliminates duplicate selected rows. This statement shows how datatype must match when columns do not exist in one or the other table:

```
SELECT part, partnum, to_date(null) date_in
   FROM orders_list1
UNION
SELECT part, to_date(null), date_in
   FROM orders_list2;
```

PART	PARTNUM	DATE_IN
SPARKPLUG	3323165	
SPARKPLUG		10/24/98
FUEL PUMP	3323162	
FUEL PUMP		12/24/99
TAILPIPE	1332999	
TAILPIPE		01/01/01
CRANKSHAFT	9394991	
CRANKSHAFT		09/12/02

```
SELECT part
   FROM orders_list1
UNION
SELECT part
   FROM orders_list2;
```

PART
SPARKPLUG
FUEL PUMP
TAILPIPE
CRANKSHAFT

**UNION ALL Example** The following statement combines the results with the UNION ALL operator, which does not eliminate duplicate selected rows:

```
SELECT part
   FROM orders_list1
UNION ALL
SELECT part
```



```
FROM orders_list2;
```

```
PART
-----
SPARKPLUG
FUEL PUMP
FUEL PUMP
TAILPIPE
CRANKSHAFT
TAILPIPE
TAILPIPE
```

Note that the `UNION` operator returns only distinct rows that appear in either result, while the `UNION ALL` operator returns all rows. A `PART` value that appears multiple times in either or both queries (such as 'FUEL PUMP') is returned only once by the `UNION` operator, but multiple times by the `UNION ALL` operator.

**INTERSECT Example** The following statement combines the results with the `INTERSECT` operator, which returns only those rows returned by both queries:

```
SELECT part
   FROM orders_list1
INTERSECT
SELECT part
   FROM orders_list2;
```

```
PART
-----
TAILPIPE
```

**MINUS Example** The following statement combines results with the `MINUS` operator, which returns only rows returned by the first query but not by the second:

```
SELECT part
   FROM orders_list1
MINUS
SELECT part
   FROM orders_list2;
```

```
PART
-----
SPARKPLUG
FUEL PUMP
```

## Other Built-In Operators

Table 3–10 lists other SQL operators.

**Table 3–10** Other SQL Operators

Operator	Purpose	Example
(+)	Indicates that the preceding column is the outer join column in a join. See <a href="#">"Outer Joins"</a> on page 5-23.	<pre>SELECT ename, dname        FROM emp, dept        WHERE dept.deptno =               emp.deptno(+);</pre>
PRIOR	Evaluates the following expression for the parent row of the current row in a hierarchical, or tree-structured, query. In such a query, you must use this operator in the <code>CONNECT BY</code> clause to define the relationship between parent and child rows. You can also use this operator in other parts of a <code>SELECT</code> statement that performs a hierarchical query. The <code>PRIOR</code> operator is a unary operator and has the same precedence as the unary <code>+</code> and <code>-</code> arithmetic operators. See <a href="#">"Hierarchical Queries"</a> on page 5-20.	<pre>SELECT empno, ename,        mgr        FROM emp        CONNECT BY               PRIOR empno =               mgr;</pre>

## User-Defined Operators

Like built-in operators, user-defined operators take a set of operands as input and return a result. However, you create them with the `CREATE OPERATOR` statement, and they are identified by names (e.g., `MERGE`). They reside in the same namespace as tables, views, types, and stand-alone functions.

Once you have defined a new operator, you can use it in SQL statements like any other built-in operator. For example, you can use user-defined operators in the select list of a `SELECT` statement, the condition of a `WHERE` clause, or in `ORDER BY` clauses and `GROUP BY` clauses. However, you must have `EXECUTE` privilege on the operator to do so, because it is a user-defined object.

For example, if you define an operator `CONTAINS`, which takes as input a text document and a keyword and returns 1 if the document contains the specified keyword, you can then write the following SQL query:

```
SELECT * FROM emp WHERE contains (resume, 'Oracle and UNIX') = 1;
```

**See Also:** ["CREATE OPERATOR"](#) on page 7-339 and *Oracle8i Data Cartridge Developer's Guide* for more information on user-defined operators.

*Form ever follows function.*

Louis Henri Sullivan, *The Tall Office Building Artistically Considered*

Functions are similar to operators in that they manipulate data items and return a result. Functions differ from operators in the format of their arguments. This format allows them to operate on zero, one, two, or more arguments:

```
function(argument, argument, ...)
```

This chapter describes two types of functions:

- [SQL Functions](#)
- [User-Defined Functions](#)

## SQL Functions


SQL functions are built into Oracle and are available for use in various appropriate SQL statements. Do not confuse SQL functions with user functions written in PL/SQL. User functions are described in ["User-Defined Functions"](#) on page 4-118. For information about functions used with Oracle interMedia, see *Oracle8i interMedia Audio, Image, and Video User's Guide and Reference*.


If you call a SQL function with an argument of a datatype other than the datatype expected by the SQL function, Oracle implicitly converts the argument to the expected datatype before performing the SQL function. See ["Data Conversion"](#) on page 2-36.

If you call a SQL function with a null argument, the SQL function automatically returns null. The only SQL functions that do not necessarily follow this behavior are `CONCAT`, `NVL`, and `REPLACE`.

In the syntax diagrams for SQL functions, arguments are indicated by their datatypes, following the conventions described in ["Syntax Diagrams and Notation"](#) in the Preface of this reference. When the parameter "function" appears in SQL syntax, replace it with one of the functions described in this section. Functions are grouped by the datatypes of their arguments and their return values. The general syntax is as follows:

**function::=**



**single\_row\_function::=**

The sections that follow list the built-in SQL functions in each of the groups illustrated above except user-defined functions. All of the built-in SQL functions are then described in alphabetical order. User-defined functions are described at the end of this chapter.

The examples provided with the function descriptions use the EMP and DEPT tables that are part of the SCOTT schema in your sample Oracle database. Many examples also use a SALES table, which has the following contents:

REGION	PRODUCT	S_DAY	S_MONTH	S_YEAR	S_AMOUNT	S_PROFIT
200	1	10	6	1998	77586	586
200	1	26	8	1998	62109	509
200	1	11	11	1998	46632	432
200	1	14	4	1999	15678	278
201	1	9	6	1998	77972	587
201	1	25	8	1998	62418	510
201	1	10	11	1998	46864	433
201	1	13	4	1999	15756	279
200	2	9	6	1998	39087	293.5
200	2	25	8	1998	31310	255
200	2	10	11	1998	23533	216.5
200	2	13	4	1999	7979	139.5
201	2	9	11	1998	23649.5	217
201	2	12	4	1999	8018.5	140
200	3	9	11	1998	15834	144.67
200	3	12	4	1999	5413.33	93.33
201	3	11	4	1999	5440	93.67
200	4	11	4	1999	4131	70.25
201	4	10	4	1999	4151.25	70.5
200	5	10	4	1999	3362	56.4
201	5	5	6	1998	16068	118.2
201	5	21	8	1998	12895.6	102.8
201	5	9	4	1999	3378.4	56.6

## Single-Row Functions

Single-row functions return a single result row for every row of a queried table or view. These functions can appear in select lists, WHERE clauses, START WITH clauses, and CONNECT BY clauses.

### Number Functions

Number functions accept numeric input and return numeric values. Most of these functions return values that are accurate to 38 decimal digits. The transcendental functions COS, COSH, EXP, LN, LOG, SIN, SINH, SQRT, TAN, and TANH are accurate to 36 decimal digits. The transcendental functions ACOS, ASIN, ATAN, and ATAN2 are accurate to 30 decimal digits. The number functions are:

---

ABS	EXP	SIGN
ACOS	FLOOR	SIN
ADD_MONTHS	LN	SINH
ATAN	LOG	SQRT
ATAN2	MOD	TAN
CEIL	POWER	TANH
COS	ROUND (Number Function)	TRUNC (Number Function)
COSH		

---

### Character Functions Returning Character Values

Character functions that return character values, unless otherwise noted, return values with the datatype VARCHAR2 and are limited in length to 4000 bytes. Functions that return values of datatype CHAR are limited in length to 2000 bytes. If the length of the return value exceeds the limit, Oracle truncates it and returns the result without an error message. The character functions that return character values are:

---

CHR	NLS_LOWER	SUBSTR
CONCAT	NLSSORT	SUBSTRB
INITCAP	NLS_UPPER	TRANSLATE
LOWER	REPLACE	TRIM
LPAD	RPAD	UPPER
LTRIM	RTRIM	
NLS_INITCAP	SOUNDEX	

---

### Character Functions Returning Number Values

The character functions that return number values are:

---

ASCII	INSTRB	LENGTHB
INSTR	LENGTH	

---

### Date Functions

Date functions operate on values of the `DATE` datatype. All date functions return a value of `DATE` datatype, except the `MONTHS_BETWEEN` function, which returns a number. The date functions are:

---

ADD_MONTHS	NEW_TIME	SYSDATE
LAST_DAY	NEXT_DAY	TRUNC (Date Function)
MONTHS_BETWEEN	ROUND (Date Function)	

---

### Conversion Functions

Conversion functions convert a value from one datatype to another. Generally, the form of the function names follows the convention *datatype TO datatype*. The first datatype is the input datatype. The second datatype is the output datatype. The SQL conversion functions are:

---

<a href="#">CHARTOROWID</a>	<a href="#">ROWIDTOCHAR</a>	<a href="#">TO_LOB</a>
<a href="#">CONVERT</a>	<a href="#">TO_CHAR</a> (date conversion)	<a href="#">TO_MULTI_BYTE</a>
<a href="#">HEXTORAW</a>	<a href="#">TO_CHAR</a> (number conversion)	<a href="#">TO_NUMBER</a>
<a href="#">NUMTODSINTERVAL</a>	<a href="#">TO_DATE</a>	<a href="#">TO_SINGLE_BYTE</a>
<a href="#">NUMTOYMINTERVAL</a>		<a href="#">TRANSLATE ... USING</a>
<a href="#">RAWTOHEX</a>		

---

### Miscellaneous Single-Row Functions

The following single-row functions do not fall into any of the other single-row function categories.

---

<a href="#">BFILENAME</a>	<a href="#">NLS_CHARSET_ID</a>	<a href="#">SYS_GUID</a>
<a href="#">DUMP</a>	<a href="#">NLS_CHARSET_NAME</a>	<a href="#">UID</a>
<a href="#">EMPTY_[B   C]LOB</a>	<a href="#">NVL</a>	<a href="#">USER</a>
<a href="#">GREATEST</a>	<a href="#">NVL2</a>	<a href="#">USERENV</a>
<a href="#">LEAST</a>	<a href="#">SYS_CONTEXT</a>	<a href="#">VSIZE</a>
<a href="#">NLS_CHARSET_DECL_LEN</a>		

---

## Aggregate Functions

Aggregate functions return a single result row based on groups of rows, rather than on single rows. Aggregate functions can appear in select lists and in `ORDER BY` and `HAVING` clauses. They are commonly used with the `GROUP BY` clause in a `SELECT` statement, where Oracle divides the rows of a queried table or view into groups. In a query containing a `GROUP BY` clause, the elements of the select list can be aggregate functions, `GROUP BY` expressions, constants, or expressions involving one of these. Oracle applies the aggregate functions to each group of rows and returns a single result row for each group.

If you omit the `GROUP BY` clause, Oracle applies aggregate functions in the select list to all the rows in the queried table or view. You use aggregate functions in the `HAVING` clause to eliminate groups from the output based on the results of the aggregate functions, rather than on the values of the individual rows of the queried table or view.

**See Also:** In the description of `SELECT`, the "[GROUP BY Examples](#)" on page 7-582 and the "[HAVING](#)" clause on page 7-579 for more information on the `GROUP BY` clause and `HAVING` clauses.



Many (but not all) aggregate functions that take a single argument accept these options:

- `DISTINCT` causes an aggregate function to consider only distinct values of the argument expression.
- `ALL` causes an aggregate function to consider all values, including all duplicates.

For example, the `DISTINCT` average of 1, 1, 1, and 3 is 2. The `ALL` average is 1.5. If you specify neither option, the default is `ALL`.

All aggregate functions except `COUNT(*)` and `GROUPING` ignore nulls. You can use the `NVL` function in the argument to an aggregate function to substitute a value for a null. `COUNT` never returns null, but returns either a number or zero. For all the remaining aggregate functions, if the data set contains no rows, or contains only rows with nulls as arguments to the aggregate function, then the function returns null.

The aggregate functions are:

<code>AVG</code>	<code>MAX</code>	<code>STDDEV_SAMP</code>
<code>CORR</code>	<code>MIN</code>	<code>SUM</code>
<code>COUNT</code>	<code>REGR_ (linear regression) functions</code>	<code>VAR_POP</code>
<code>COVAR_POP</code>	<code>STDDEV</code>	<code>VAR_SAMP</code>
<code>COVAR_SAMP</code>	<code>STDDEV_POP</code>	<code>VARIANCE</code>
<code>GROUPING</code>		


## Analytic Functions

Analytic functions compute an aggregate value based on a group of rows. The group of rows is called a **window** and is defined by the analytic clause. For each row, a "sliding" window of rows is defined. The window determines the range of rows used to perform the calculations for the "current row". Window sizes can be based on either a physical number of rows or a logical interval such as time.


Analytic functions are the last set of operations performed in a query except for the final `ORDER BY` clause. All joins and all `WHERE`, `GROUP BY`, and `HAVING` clauses are completed before the analytic functions are processed. Therefore, analytic functions can appear only in the select list or `ORDER BY` clause.

Analytic functions are commonly used to compute cumulative, moving, centered, and reporting aggregates.


**analytic\_function::=**




**analytic\_clause::=**



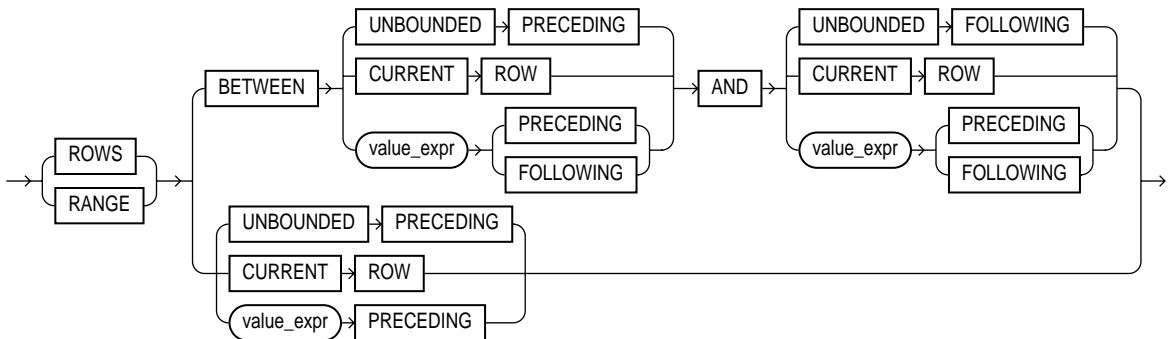
**query\_partition\_clause::=**



**ORDER\_BY\_clause::=**



**windowing\_clause::=**



The keywords and parameters of this syntax are:

---

<i>analytic_function</i>	is the name of an analytic function (see the listings of different types of analytic functions following this table).
<i>arguments</i>	Analytic functions take 0 to 3 arguments.
OVER	indicates that the function operates on a query result set. That is, it is computed after the FROM, WHERE, GROUP BY, and HAVING clauses. You can specify analytic functions with this clause in the select list or ORDER BY clause. To filter the results of a query based on an analytic function, nest these functions within the parent query, and then filter the results of the nested subquery.
<i>analytic_clause</i>	
<i>query_partition_clause</i>	
PARTITION BY	partitions the query result set into groups based on one or more <i>value_expr</i> . If you omit this clause, the function treats all rows of the query result set as a single group.  You can specify multiple analytic functions in the same query, each with the same or different PARTITION BY keys.
	<b>Note:</b> If the objects being queried have the parallel attribute, and if you specify an analytic function with the <i>query_partition_clause</i> , then the function computations are parallelized as well.
<i>value_expr</i>	Valid value expressions are constants, columns, nonanalytic functions, function expressions, or expressions involving any of these.
<i>ORDER_BY_clause</i>	
ORDER BY	specifies how data is ordered within a partition. You can order the values in a partition on multiple keys, each defined by a <i>value_expr</i> and each qualified by an ordering sequence.  Within each function, you can specify multiple ordering expressions. Doing so is especially useful when using functions that rank values, because the second expression can resolve ties between identical values for the first expression.  <b>See Also:</b> <i>order_by_clause</i> of "SELECT and Subqueries" on page 7-580 for more information on this clause.
	<b>Note:</b> Analytic functions always operate on rows in the order specified in the <i>ORDER_BY_clause</i> of the function. However, the <i>ORDER_BY_clause</i> of the function does not guarantee the order of the result. Use the <i>ORDER_BY_clause</i> of the query to guarantee the final result ordering.
	<b>Restriction:</b> When used in an analytic function, the <i>ORDER_BY_clause</i> must take an expression ( <i>expr</i> ). Position ( <i>position</i> ) and column aliases ( <i>c_alias</i> ) are invalid. Otherwise this <i>ORDER_BY_clause</i> is the same as that used to order the overall query or subquery.

---

---

ASC   DESC	specifies the ordering sequence (ascending or descending). ASC is the default.
NULLS FIRST   NULLS LAST	specifies whether returned rows containing null values should appear first or last in the ordering sequence.  NULLS LAST is the default for ascending order, and NULLS FIRST is the default for descending order.

**windowing\_clause**

ROWS   RANGE	<p>These keywords define for each row a "window" (a physical or logical set of rows) used for calculating the function result. The function is then applied to all the rows in the window. The window "slides" through the query result set or partition from top to bottom.</p> <ul style="list-style-type: none"> <li>■ ROWS specifies the window in physical units (rows).</li> <li>■ RANGE specifies the window as a logical offset.</li> </ul> <p>You cannot specify this clause unless you have specified the <i>ORDER_BY_clause</i>.</p>
--------------	---

---

**Note:** The value returned by an analytic function with a logical offset is always deterministic. However, the value returned by an analytic function with a physical offset may produce nondeterministic results unless the ordering expression results in a unique ordering. You may have to specify multiple columns in the *ORDER\_BY\_clause* to achieve this unique ordering.

---

BETWEEN . . . AND	<p>lets you specify a start point and end point for the window. The first expression (before AND) defines the start point and the second expression (after AND) defines the end point.</p> <p>If you omit BETWEEN and specify only one end point, Oracle considers it the start point, and the end point defaults to the current row.</p>
UNBOUNDED PRECEDING	specifies that the window starts at the first row of the partition. This is the start point specification and cannot be used as an end point specification.
UNBOUNDED FOLLOWING	specifies that the window ends at the last row of the partition. This is the end point specification and cannot be used as a start point specification.
CURRENT ROW	<p>As a start point, CURRENT ROW specifies that the window begins at the current row or value (depending on whether you have specified ROW or RANGE, respectively). In this case the end point cannot be <i>value_expr</i> PRECEDING.</p> <p>As an end point, CURRENT ROW specifies that the window ends at the current row or value (depending on whether you have specified ROW or RANGE, respectively). In this case the start point cannot be <i>value_expr</i> FOLLOWING.</p>

---

<i>value_expr</i> PRECEDING	For RANGE or ROW:
<i>value_expr</i> FOLLOWING	<ul style="list-style-type: none"> <li>■ If <i>value_expr</i> FOLLOWING is the start point, then the end point must be <i>value_expr</i> FOLLOWING.</li> <li>■ If <i>value_expr</i> PRECEDING is the end point, then the start point must be <i>value_expr</i> PRECEDING.</li> </ul>

If you are defining a logical window defined by an interval of time in numeric format, you may need to use conversion functions.

**See Also:** "[NUMTOYMINTERVAL](#)" on page 4-63 and "[NUMTODSINTERVAL](#)" on page 4-62 for information on converting numeric times into interval literals.

If you specified ROWS:

- *value\_expr* is a physical offset. It must be a constant or expression and must evaluate to a positive numeric value.
- If *value\_expr* is part of the start point, it must evaluate to a row before the end point.

If you specified RANGE:

- *value\_expr* is a logical offset. It must be a constant or expression that evaluates to a positive numeric value or an interval literal.

**See Also:** "[Literals](#)" on page 2-2 for information on interval literals.

- You can specify only one expression in the *ORDER\_BY\_clause*
- If *value\_expr* evaluates to a numeric value, the ORDER BY *expr* must be a NUMBER or DATE datatype.
- If *value\_expr* evaluates to an interval value, the ORDER BY *expr* must be a DATE datatype.

If you omit the *windowing\_clause* entirely, the default is RANGE BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW.

---

Analytic functions are commonly used in data warehousing environments. The analytic functions are:

AVG	LEAD	STDDEV
CORR	MAX	STDDEV_POP
COVAR_POP	MIN	STDDEV_SAMP
COVAR_SAMP	NTILE	SUM
COUNT	PERCENT_RANK	VAR_POP
CUME_DIST	RATIO_TO_REPORT	VAR_SAMP
DENSE_RANK	RANK	VARIANCE
LAG	REGR_ (linear regression) functions	
FIRST_VALUE		
LAST_VALUE	ROW_NUMBER	

**See Also:** *Oracle8i Data Warehousing Guide* for more information on these functions, and for scenarios illustrating their use.

## Object Reference Functions

Object functions manipulate REFs, which are references to objects of specified object types. The object reference functions are:


DEREF	REF	VALUE
MAKE_REF	REFTOHEX	

**See Also:** *Oracle8i Concepts* and *Oracle8i Application Developer's Guide - Fundamentals* for more information about REFs.

## Alphabetical Listing of SQL Functions

### ABS

#### Syntax



**Purpose**

ABS returns the absolute value of *n*.

**Example**

```
SELECT ABS(-15) "Absolute" FROM DUAL;
```

```

Absolute
-----
          15

```

## ACOS

**Syntax**

```
→ [ACOS] ( ( ) n ( ) ) →
```

**Purpose**

ACOS returns the arc cosine of *n*. Inputs are in the range of -1 to 1, and outputs are in the range of 0 to  $\pi$  and are expressed in radians.

**Example**

```
SELECT ACOS(.3) "Arc_Cosine" FROM DUAL;
```

```

Arc_Cosine
-----
1.26610367

```

## ADD\_MONTHS

**Syntax**

```
→ [ADD_MONTHS] ( ( ) d , ( ) n ( ) ) →
```

**Purpose**

ADD\_MONTHS returns the date *d* plus *n* months. The argument *n* can be any integer. If *d* is the last day of the month or if the resulting month has fewer days than the

day component of *d*, then the result is the last day of the resulting month. Otherwise, the result has the same day component as *d*.

### Example

```
SELECT TO_CHAR(  
    ADD_MONTHS(hiredate,1),  
    'DD-MON-YYYY') "Next month"  
FROM emp  
WHERE ename = 'SMITH';
```

```
Next Month  
-----  
17-JAN-1981
```

## ASCII

### Syntax

→ ASCII ( char ) →

### Purpose

ASCII returns the decimal representation in the database character set of the first character of *char*. If your database character set is 7-bit ASCII, this function returns an ASCII value. If your database character set is EBCDIC Code Page 500, this function returns an EBCDIC value. There is no similar EBCDIC character function.

### Example

```
SELECT ASCII('Q') FROM DUAL;
```

```
ASCII('Q')  
-----  
81
```



# ASIN

## Syntax



## Purpose

ASIN returns the arc sine of  $n$ . Inputs are in the range of -1 to 1, and outputs are in the range of  $-\pi/2$  to  $\pi/2$  and are expressed in radians.

## Example

```
SELECT ASIN(.3) "Arc_Sine" FROM DUAL;
```

```

Arc_Sine
-----
.304692654

```

# ATAN

## Syntax



## Purpose

ATAN returns the arc tangent of  $n$ . Inputs are in an unbounded range, and outputs are in the range of  $-\pi/2$  to  $\pi/2$  and are expressed in radians.

## Example

```
SELECT ATAN(.3) "Arc_Tangent" FROM DUAL;
```


```

Arc_Tangent
-----
.291456794

```

# ATAN2

## Syntax



## Purpose

ATAN2 returns the arc tangent of *n* and *m*. Inputs are in an unbounded range, and outputs are in the range of  $-\pi$  to  $\pi$ , depending on the signs of *n* and *m*, and are expressed in radians. ATAN2(*n*,*m*) is the same as ATAN2(*n*/*m*)


## Example

```
SELECT ATAN2(.3, .2) "Arc_Tangent2" FROM DUAL;
```

```
Arc_Tangent2
-----
.982793723
```

# AVG

## Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

## Purpose

AVG returns average value of *expr*. You can use it as an aggregate or analytic function.

If you specify **DISTINCT**, you can specify only the *query\_partition\_clause* of the *analytic\_clause*. The *ORDER\_BY\_clause* and *windowing\_clause* are not allowed.

**See Also:** "[Aggregate Functions](#)" on page 4-6.

## Aggregate Example

The following example calculates the average salary of all employees in the EMP table:

```
SELECT AVG(sal) "Average" FROM emp;
```

```

Average
-----
2077.21429
```

## Analytic Example

The following example calculates, for each employee in the EMP table, the average salary of the employees reporting to the same manager who were hired in the range just before through just after the employee:

```
SELECT mgr, ename, hiredate, sal,
       AVG(sal) OVER (PARTITION BY mgr ORDER BY hiredate
                     ROWS BETWEEN 1 PRECEDING AND 1 FOLLOWING) AS c_mavg
FROM emp;
```

MGR	ENAME	HIREDATE	SAL	C_MAVG
7566	FORD	03-DEC-81	3000	3000
7566	SCOTT	19-APR-87	3000	3000
7698	ALLEN	20-FEB-81	1600	1425
7698	WARD	22-FEB-81	1250	1450
7698	TURNER	08-SEP-81	1500	1333.33333
7698	MARTIN	28-SEP-81	1250	1233.33333
7698	JAMES	03-DEC-81	950	1100
7782	MILLER	23-JAN-82	1300	1300
7788	ADAMS	23-MAY-87	1100	1100
7839	JONES	02-APR-81	2975	2912.5
7839	BLAKE	01-MAY-81	2850	2758.33333
7839	CLARK	09-JUN-81	2450	2650
7902	SMITH	17-DEC-80	800	800
	KING	17-NOV-81	5000	5000

## BFILENAME

### Syntax

```
→ BFILENAME ( '(' directory ')' , filename ) →
```

## Purpose

`BFILENAME` returns a `BFILE` locator that is associated with a physical LOB binary file on the server's file system. A directory is an alias for a full pathname on the server's file system where the files are actually located, and 'filename' is the name of the file in the server's file system.

Neither 'directory' nor 'filename' needs to point to an existing object on the file system at the time you specify `BFILENAME`. However, you must associate a `BFILE` value with a physical file before performing subsequent SQL, PL/SQL, `DBMS_LOB` package, or OCI operations.

### See Also:

- *Oracle8i Application Developer's Guide - Large Objects (LOBs) and Oracle Call Interface Programmer's Guide* for more information on LOBs.
- ["CREATE DIRECTORY"](#) on page 7-282.

## Example

```
INSERT INTO file_tbl
VALUES (BFILENAME ('lob_dir1', 'image1.gif'));
```

# CEIL

## Syntax

→ `CEIL` → ( → `n` → ) →

## Purpose

`CEIL` returns smallest integer greater than or equal to *n*.

## Example


```
SELECT CEIL(15.7) "Ceiling" FROM DUAL;
```

```

Ceiling
-----
      16
```

## CHARTOROWID

### Syntax



### Purpose

CHARTOROWID converts a value from CHAR or VARCHAR2 datatype to ROWID datatype.


### Example

```
SELECT ename FROM emp
       WHERE ROWID = CHARTOROWID('AAAafZAABAAACp8AAO');
```

```
ENAME
-----
LEWIS
```

## CHR

### Syntax



### Purpose

CHR returns the character having the binary equivalent to *n* in either the database character set or the national character set.

If USING NCHAR\_CS is *not* specified, this function returns the character having the binary equivalent to *n* as a VARCHAR2 value in the database character set.

If USING NCHAR\_CS is specified, this function returns the character having the binary equivalent to *n* as a NVARCHAR2 value in the national character set.

### Example 1

```
SELECT CHR(67) || CHR(65) || CHR(84) "Dog"
       FROM DUAL;
```

```
Dog
---
CAT
```


### Example 2

```
SELECT CHR(16705 USING NCHAR_CS) FROM DUAL;
```

```
C
-
A
```

## CONCAT

### Syntax



### Purpose

CONCAT returns *char1* concatenated with *char2*. This function is equivalent to the concatenation operator (||). For information on this operator, see "[Concatenation Operator](#)" on page 3-3.

### Example


This example uses nesting to concatenate three character strings:

```
SELECT CONCAT(CONCAT(ename, ' is a '), job) "Job"
FROM emp
WHERE empno = 7900;
```

```
Job
-----
JAMES is a CLERK
```

# CONVERT

## Syntax



## Purpose

CONVERT returns smallest integer greater than or equal to *expr*.

Converts a character string from one character set to another.

- The *char* argument is the value to be converted.
- The *dest\_char\_set* argument is the name of the character set to which *char* is converted.
- The *source\_char\_set* argument is the name of the character set in which *char* is stored in the database. The default value is the database character set.

Both the destination and source character set arguments can be either literals or columns containing the name of the character set.

For complete correspondence in character conversion, it is essential that the destination character set contains a representation of all the characters defined in the source character set. Where a character does not exist in the destination character set, a replacement character appears. Replacement characters can be defined as part of a character set definition.

## Example

```
SELECT CONVERT('Groß', 'US7ASCII', 'WE8HP')
       "Conversion" FROM DUAL;
```

```
Conversion
-----
Gross
```


Common character sets include:

- US7ASCII: US 7-bit ASCII character set
- WE8DECDEC: West European 8-bit character set

- WE8HP: HP West European Laserjet 8-bit character set
- F7DEC: DEC French 7-bit character set
- WE8EBCDIC500: IBM West European EBCDIC Code Page 500
- WE8PC850: IBM PC Code Page 850
- WE8ISO8859P1: ISO 8859-1 West European 8-bit character set

## CORR

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

CORR returns the coefficient of correlation of a set of number pairs. You can use it as an aggregate or analytic function.

Both *expr1* and *expr2* are number expressions. Oracle applies the function to the set of (*expr1* , *expr2*) after eliminating the pairs for which either *expr1* or *expr2* is null. Then Oracle makes the following computation:

$$\text{COVAR\_POP}(\text{expr1}, \text{expr2}) / (\text{STDDEV\_POP}(\text{expr1}) * \text{STDDEV\_POP}(\text{expr2}))$$

The function returns a value of type NUMBER. If the function is applied to an empty set, it returns null.

**See Also:** "[Aggregate Functions](#)" on page 4-6.

### Aggregate Example

The following example calculates the coefficient of correlation between the salaries and commissions of the employees whose manager is 7698 from the EMP table:

```

SELECT mgr, CORR(sal, comm) FROM EMP
   GROUP BY mgr
   HAVING mgr = 7698;

      MGR CORR(SAL,COMM)

```



```
-----
          7698      -.69920974
```

### Analytic Example

The following example returns the cumulative coefficient of correlation of monthly sales and monthly profits from the SALES table for year 1998:

```
SELECT s_month, CORR(SUM(s_amount), SUM(s_profit))
       OVER (ORDER BY s_month) AS CUM_CORR
FROM sales
WHERE s_year=1998
GROUP BY s_month
ORDER BY s_month;
```

```
S_MONTH      CUM_CORR
-----
          6
          8          1
         11 .860554259
```

Correlation functions require more than one row on which to operate, so the first row in the preceding example has no value calculated for it.

## COS

### Syntax



### Purpose

COS returns the cosine of  $n$  (an angle expressed in radians).

### Example

```
SELECT COS(180 * 3.14159265359/180)
"Cosine of 180 degrees" FROM DUAL;
```

```
Cosine of 180 degrees
-----
                    -1
```

## COSH

### Syntax



### Purpose

COSH returns the hyperbolic cosine of *n*.


### Example

```
SELECT COSH(0) "Hyperbolic cosine of 0" FROM DUAL;
```

```
Hyperbolic cosine of 0
-----
1
```

## COUNT

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

COUNT returns the number of rows in the query. You can use it as an aggregate or analytic function.

If you specify *DISTINCT*, you can specify only the *query\_partition\_clause* of the *analytic\_clause*. The *ORDER\_BY\_clause* and *windowing\_clause* are not allowed.

If you specify *expr*, COUNT returns the number of rows where *expr* is not null. You can count either all rows, or only distinct values of *expr*.

If you specify the asterisk (\*), this function returns all rows, including duplicates and nulls. COUNT never returns null.

**See Also:** ["Aggregate Functions"](#) on page 4-6.

### Aggregate Examples

```
SELECT COUNT(*) "Total" FROM emp;
```

```
Total
-----
      14
```

```
SELECT COUNT(mgr) "Count" FROM emp;
```

```
Count
-----
      13
```

```
SELECT COUNT(DISTINCT mgr) "Managers" FROM emp;
```

```
Managers
-----
         6
```

### Analytic Example

The following example calculates, for each employee in the EMP table, the moving count of employees earning salaries in the range \$50 less than through \$150 greater than the employee's salary.


```
SELECT ename, sal,
       COUNT(*) OVER (ORDER BY sal RANGE BETWEEN 50 PRECEDING
                      AND 150 FOLLOWING) AS mov_count
FROM emp;
```

ENAME	SAL	MOV_COUNT
SMITH	800	2
JAMES	950	2
ADAMS	1100	3
WARD	1250	3
MARTIN	1250	3
MILLER	1300	3
TURNER	1500	2
ALLEN	1600	1

CLARK	2450	1
BLAKE	2850	4
JONES	2975	3
SCOTT	3000	3
FORD	3000	3
KING	5000	1

## COVAR\_POP

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

### Purpose

COVAR\_POP returns the population covariance of a set of number pairs. You can use it as an aggregate or analytic function.

Both *expr1* and *expr2* are number expressions. Oracle applies the function to the set of (*expr1* , *expr2*) pairs after eliminating all pairs for which either *expr1* or *expr2* is null. Then Oracle makes the following computation:

$$(\text{SUM}(\text{expr1} * \text{expr2}) - \text{SUM}(\text{expr2}) * \text{SUM}(\text{expr1}) / n) / n$$

where *n* is the number of (*expr1* , *expr2*) pairs where neither *expr1* nor *expr2* is null.

The function returns a value of type NUMBER. If the function is applied to an empty set, it returns null.

**See Also:** ["Aggregate Functions"](#) on page 4-6.

### Aggregate Example

The following example calculates the population covariance for the amount of sales and sale profits for each year from the table SALES.

```
SELECT s_year,
       COVAR_POP(s_amount, s_profit) AS COVAR_POP,
       COVAR_SAMP(s_amount, s_profit) AS COVAR_SAMP
FROM sales GROUP BY s_year;
```


S_YEAR	COVAR_POP	COVAR_SAMP
1998	3747965.53	4060295.99
1999	360536.162	400595.736

### Analytic Example

See "[COVAR\\_SAMP](#)" on page 4-27.

## COVAR\_SAMP

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

COVAR\_SAMP returns the sample covariance of a set of number pairs. You can use it as an aggregate or analytic function.

Both *expr1* and *expr2* are number expressions. Oracle applies the function to the set of (*expr1*, *expr2*) pairs after eliminating all pairs for which either *expr1* or *expr2* is null. Then Oracle makes the following computation:

$$(\text{SUM}(\text{expr1} * \text{expr2}) - \text{SUM}(\text{expr1}) * \text{SUM}(\text{expr2}) / n) / (n-1)$$

where *n* is the number of (*expr1*, *expr2*) pairs where neither *expr1* nor *expr2* is null.

The function returns a value of type NUMBER. If the function is applied to an empty set, it returns null.

**See Also:** "[Aggregate Functions](#)" on page 4-6.

### Aggregate Example

See "[COVAR\\_POP](#)" on page 4-26.

### Analytic Example


The following example calculates cumulative sample covariance of the amount of sales and sale profits in 1998.

```
SELECT s_year, s_month, s_day,
       COVAR_POP(s_amount, s_profit)
         OVER (ORDER BY s_month, s_day) AS CUM_COVP,
       COVAR_SAMP(s_amount, s_profit)
         OVER (ORDER BY s_month, s_day) AS CUM_COVS
FROM sales
WHERE s_year=1998
ORDER BY s_year, s_month, s_day;
```

S_YEAR	S_MONTH	S_DAY	CUM_COVP	CUM_COVS
1998	6	5	5	0
1998	6	9	9	4940952.6
1998	6	9	9	4940952.6
1998	6	10	10	5281752.33
1998	8	21	21	6092799.46
1998	8	25	25	4938283.61
1998	8	25	25	4938283.61
1998	8	26	26	4612074.09
1998	11	9	9	4556799.53
1998	11	9	9	4556799.53
1998	11	10	10	4014833.65
1998	11	10	10	4014833.65
1998	11	11	11	3747965.53

## CUME\_DIST

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

CUME\_DIST (cumulative distribution) is an analytic function. It computes the relative position of a specified value in a group of values. For a row R, assuming

ascending ordering, the `CUME_DIST` of `R` is the number of rows with values lower than or equal to the value of `R`, divided by the number of rows being evaluated (the entire query result set or a partition). The range of values returned by `CUME_DIST` is  $>0$  to  $\leq 1$ . Tie values always evaluate to the same cumulative distribution value.

### Example


The following example calculates the salary percentile for each employee within each job category excluding job categories `PRESIDENT` and `MANAGER`. For example, 50% of clerks have salaries less than or equal to James.

```
SELECT job, ename, sal, CUME_DIST()
      OVER (PARTITION BY job ORDER BY sal) AS cume_dist
FROM emp
WHERE job NOT IN ('MANAGER', 'PRESIDENT');
```

JOB	ENAME	SAL	CUME_DIST
ANALYST	SCOTT	3000	1
ANALYST	FORD	3000	1
CLERK	SMITH	800	.25
CLERK	JAMES	950	.5
CLERK	ADAMS	1100	.75
CLERK	MILLER	1300	1
SALESMAN	WARD	1250	.5
SALESMAN	MARTIN	1250	.5
SALESMAN	TURNER	1500	.75
SALESMAN	ALLEN	1600	1

## DENSE\_RANK

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

`DENSE_RANK` is an analytic function. It computes the rank of each row returned from a query with respect to the other rows, based on the values of the *value\_exprs*

in the *ORDER\_BY\_clause*. Rows with equal values for the ranking criteria receive the same rank. The ranks are consecutive integers beginning with 1. The largest rank value is the number of unique values returned by the query. Rank values are not skipped in the event of ties.

### Example

The following statement selects the department name, employee name, and salary of all employees who work in the RESEARCH or SALES department, and then computes a rank for each unique salary in each of the two departments. The salaries that are equal receive the same rank. Compare this example with the example for "RANK" on page 4-67.

```
SELECT dname, ename, sal, DENSE_RANK()
   OVER (PARTITION BY dname ORDER BY sal) as drank
   FROM emp, dept
   WHERE emp.deptno = dept.deptno
   AND dname IN ('SALES', 'RESEARCH');
```

DNAME	ENAME	SAL	DRANK
RESEARCH	SMITH	800	1
RESEARCH	ADAMS	1100	2
RESEARCH	JONES	2975	3
RESEARCH	FORD	3000	4
RESEARCH	SCOTT	3000	4
SALES	JAMES	950	1
SALES	MARTIN	1250	2
SALES	WARD	1250	2
SALES	TURNER	1500	3
SALES	ALLEN	1600	4
SALES	BLAKE	2850	5

## DEREF

### Syntax





## Purpose

DEREF returns the object reference of argument *expr*, where *expr* must return a REF to an object.

## Example


```
CREATE TYPE emp_type AS OBJECT
  (eno NUMBER, ename VARCHAR2(20), salary NUMBER);
CREATE TABLE emp_table OF emp_type
  (primary key (eno, ename));
CREATE TABLE dept_table
  (dno NUMBER, mgr REF emp_type SCOPE IS emp_table);
INSERT INTO emp_table VALUES (10, 'jack', 50000);
INSERT INTO dept_table SELECT 10, REF(e) FROM emp_table e;
SELECT Deref(mgr) from dept_table;
```

```
DEREF(MGR)(ENO, ENAME, SALARY)
```

```
-----
EMP_TYPE(10, 'jack', 50000)
```

# DUMP

## Syntax



## Purpose

DUMP returns a VARCHAR2 value containing the datatype code, length in bytes, and internal representation of *expr*. The returned result is always in the database character set. For the datatype corresponding to each code, see [Table 2-1](#) on page 2-13.

The argument *return\_fmt* specifies the format of the return value and can have any of the following values:

- 8 returns result in octal notation
- 10 returns result in decimal notation

- 16 returns result in hexadecimal notation
- 17 returns result as single characters

By default, the return value contains no character set information. To retrieve the character set name of *expr*, specify any of the format values above, plus 1000. For example, a *return\_fmt* of 1008 returns the result in octal, plus provides the character set name of *expr*.

The arguments *start\_position* and *length* combine to determine which portion of the internal representation to return. The default is to return the entire internal representation in decimal notation.

If *expr* is null, this function returns a null.

### Example 1

```
SELECT DUMP('abc', 1016)
       FROM DUAL;
```

```
DUMP('ABC',1016)
-----
Type=96 Len=3 CharacterSet=WE8DEC: 61,62,63
```

### Example 2

```
SELECT DUMP(ename, 8, 3, 2) "OCTAL"
       FROM emp
       WHERE ename = 'SCOTT';
```

```
OCTAL
-----
Type=1 Len=5: 117,124
```


### Example 3

```
SELECT DUMP(ename, 10, 3, 2) "ASCII"
       FROM emp
       WHERE ename = 'SCOTT';
```

```
ASCII
-----
Type=1 Len=5: 79,84
```

## EMPTY\_[B | C]LOB

### Syntax



### Purpose

EMPTY\_BLOB and EMPTY\_CLOB returns an empty LOB locator that can be used to initialize a LOB variable or in an INSERT or UPDATE statement to initialize a LOB column or attribute to EMPTY. EMPTY means that the LOB is initialized, but not populated with data.

You cannot use the locator returned from this function as a parameter to the DBMS\_LOB package or the OCI.

### Example


```

INSERT INTO lob_tab1 VALUES (EMPTY_BLOB());
UPDATE lob_tab1
   SET clob_col = EMPTY_BLOB();

```

## EXP

### Syntax



### Purpose

EXP returns e raised to the *n*th power, where e = 2.71828183 ...

### Example

```

SELECT EXP(4) "e to the 4th power" FROM DUAL;

```

```

e to the 4th power
-----
          54.59815

```

## FIRST\_VALUE

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

`FIRST_VALUE` is an analytic function. It returns the first value in an ordered set of values.

### Examples

The following example selects, for each employee in Department 20, the name of the employee with the highest salary.

```

SELECT deptno, ename, sal, FIRST_VALUE(ename)
      OVER (ORDER BY sal DESC ROWS UNBOUNDED PRECEDING) AS rich_emp
FROM (SELECT * FROM emp WHERE deptno = 20 ORDER BY empno);

```

DEPTNO	ENAME	SAL	RICH_EMP
20	SCOTT	3000	SCOTT
20	FORD	3000	SCOTT
20	JONES	2975	SCOTT
20	ADAMS	1100	SCOTT
20	SMITH	800	SCOTT

The example illustrates the nondeterministic nature of the `FIRST_VALUE` function. Scott and Ford have the same salary, so are in adjacent rows. Scott appears first because the rows returned by the subquery are ordered by `EMPNO`. However, if the rows returned by the subquery are ordered by `EMPNO` in descending order, as in the next example, the function returns a different value:

```

SELECT deptno, ename, sal, FIRST_VALUE(ename)
      OVER (ORDER BY sal DESC ROWS UNBOUNDED PRECEDING) AS fv
FROM (SELECT * FROM emp WHERE deptno = 20 ORDER BY empno desc);

```

DEPTNO	ENAME	SAL	FV
20	FORD	3000	FORD
20	SCOTT	3000	FORD

```

20 JONES          2975 FORD
20 ADAMS          1100 FORD
20 SMITH          800  FORD
    
```

The following example shows how to make the FIRST\_VALUE function deterministic by ordering on a unique key.


```

SELECT deptno, ename, sal, hiredate, FIRST_VALUE(ename)
       OVER (ORDER BY sal DESC, hiredate ROWS UNBOUNDED PRECEDING) AS fv
FROM (SELECT * FROM emp WHERE deptno = 20 ORDER BY empno desc);
    
```

DEPTNO	ENAME	SAL	HIREDATE	FV
20	FORD	3000	03-DEC-81	FORD
20	SCOTT	3000	19-APR-87	FORD
20	JONES	2975	02-APR-81	FORD
20	ADAMS	1100	23-MAY-87	FORD
20	SMITH	800	17-DEC-80	FORD

## FLOOR

### Syntax



### Purpose

FLOOR returns largest integer equal to or less than *n*.

### Example

```


SELECT FLOOR(15.7) "Floor" FROM DUAL;
    
```

```

      Floor
-----
      15
    
```

## GREATEST

### Syntax



### Purpose

**GREATEST** returns the greatest of the list of *exprs*. All *exprs* after the first are implicitly converted to the datatype of the first *expr* before the comparison. Oracle compares the *exprs* using nonpadded comparison semantics. Character comparison is based on the value of the character in the database character set. One character is greater than another if it has a higher character set value. If the value returned by this function is character data, its datatype is always `VARCHAR2`.


### Example

```
SELECT GREATEST ('HARRY', 'HARRIOT', 'HAROLD')
       "Greatest" FROM DUAL;
```

```
Greatest
-----
HARRY
```

## GROUPING

### Syntax



### Purpose

This function is applicable only in a `SELECT` statement that contains a `GROUP BY` extension, such as `ROLLUP` or `CUBE`. These operations produce superaggregate rows that contain nulls representing the set of all values. You can use the `GROUPING` function to distinguish a null that represents the set of all values in a superaggregate row from an actual null.

The *expr* in the GROUPING function must match one of the expressions in the GROUP BY clause. The function returns a value of 1 if the value of *expr* in the row is a null representing the set of all values. Otherwise, it returns zero. The datatype of the value returned by the GROUPING function is Oracle NUMBER. See the [group\\_by\\_clause](#) of the SELECT statement on page 7-578 for a discussion of these terms.

### Example

In the following example, if the GROUPING function returns 1 (indicating a superaggregate row rather than a data row from the table), the string "All Jobs" appears instead of the null that would otherwise appear:

```
SELECT DECODE(GROUPING(dname), 1, 'All Departments',
             dname) AS dname,
       DECODE(GROUPING(job), 1, 'All Jobs', job) AS job,
       COUNT(*) "Total Empl", AVG(sal) * 12 "Average Sal"
FROM emp, dept
WHERE dept.deptno = emp.deptno
GROUP BY ROLLUP (dname, job);
```

DNAME	JOB	Total Empl	Average Sa
ACCOUNTING	CLERK	1	15600
ACCOUNTING	MANAGER	1	29400
ACCOUNTING	PRESIDENT	1	60000
ACCOUNTING	All Jobs	3	35000
RESEARCH	ANALYST	2	36000
RESEARCH	CLERK	2	11400
RESEARCH	MANAGER	1	35700
RESEARCH	All Jobs	5	26100
SALES	CLERK	1	11400
SALES	MANAGER	1	34200
SALES	SALESMAN	4	16800
SALES	All Jobs	6	18800
All Departments	All Jobs	14	24878.5714

## HEXTORAW

### Syntax

→ HEXTORAW → ( ( char ) ) →

**Purpose**

HEXTORAW converts *char* containing hexadecimal digits to a raw value.

**Example**

```
INSERT INTO graphics (raw_column)
  SELECT HEXTORAW('7D') FROM DUAL;
```

## INITCAP

**Syntax**
**Purpose**

INITCAP returns *char*, with the first letter of each word in uppercase, all other letters in lowercase. Words are delimited by white space or characters that are not alphanumeric.

**Example**

```
SELECT INITCAP('the soap') "Capitals" FROM DUAL;
```

```
Capitals
-----
The Soap
```

## INSTR

**Syntax**
**Purpose**

INSTR searches *char1* beginning with its *n*th character for the *m*th occurrence of *char2* and returns the position of the character in *char1* that is the first character of



this occurrence. If *n* is negative, Oracle counts and searches backward from the end of *char1*. The value of *m* must be positive. The default values of both *n* and *m* are 1, meaning Oracle begins searching at the first character of *char1* for the first occurrence of *char2*. The return value is relative to the beginning of *char1*, regardless of the value of *n*, and is expressed in characters. If the search is unsuccessful (if *char2* does not appear *m* times after the *n*th character of *char1*) the return value is 0.

### Example 1

```
SELECT INSTR('CORPORATE FLOOR','OR', 3, 2)
       "Instring" FROM DUAL;
```

```
    Instring
-----
           14
```


### Example 2

```
SELECT INSTR('CORPORATE FLOOR','OR', -3, 2)
       "Reversed Instring"
       FROM DUAL;
```

```
Reversed Instring
-----
                2
```

## INSTRB

### Syntax



### Purpose

INSTRB is the same as INSTR, except that *n* and the return value are expressed in bytes, rather than in characters. For a single-byte database character set, INSTRB is equivalent to INSTR.

## Example


This example assumes a double-byte database character set.

```
SELECT INSTRB('CORPORATE FLOOR', 'OR', 5, 2)
       "Instring in bytes"
FROM DUAL;
```

```
Instring in bytes
-----
                      27
```

## LAG

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

LAG is an analytic function. It provides access to more than one row of a table at the same time without a self-join. Given a series of rows returned from a query and a position of the cursor, LAG provides access to a row at a given physical offset prior to that position.

If you do not specify *offset*, its default is 1. The optional *default* value is returned if the offset goes beyond the scope of the window. If you do not specify *default*, its default value is null.

### Example

The following example provides, for each salesperson in the EMP table, the salary of the employee hired just before:

```
SELECT ename, hiredate, sal,
       LAG(sal, 1, 0) OVER (ORDER BY hiredate) as prev_sal
FROM emp
WHERE job = 'SALESMAN';
```

```
ENAME          HIREDATE          SAL    PREV_SAL
```

```

-----
ALLEN      20-FEB-81      1600          0
WARD       22-FEB-81      1250         1600
TURNER     08-SEP-81      1500         1250
MARTIN     28-SEP-81      1250         1500

```

## LAST\_DAY

### Syntax

```

→ LAST_DAY → ( → ( → d → ) → ) →

```

### Purpose

LAST\_DAY returns the date of the last day of the month that contains *d*. You might use this function to determine how many days are left in the current month.

### Example 1

```

SELECT SYSDATE,
       LAST_DAY(SYSDATE) "Last",
       LAST_DAY(SYSDATE) - SYSDATE "Days Left"
FROM DUAL;

```

```

SYSDATE   Last           Days Left
-----
23-OCT-97 31-OCT-97          8

```

### Example 2

```

SELECT TO_CHAR(
       ADD_MONTHS(
         LAST_DAY(hiredate), 5),
       'DD-MON-YYYY') "Five months"
FROM emp
WHERE ename = 'MARTIN';

```


```

Five months
-----
28-FEB-1982

```

## LAST\_VALUE

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

LAST\_VALUE is an analytic function. It returns the last value in an ordered set of values.

### Examples

The following example returns the hiredate of the employee earning the highest salary.

```

SELECT ename, sal, hiredate, LAST_VALUE(hiredate) OVER
  (ORDER BY sal
   ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING) AS lv
FROM (SELECT * FROM emp WHERE deptno=20 ORDER BY hiredate);

```

ENAME	SAL	HIREDATE	LV
SMITH	800	17-DEC-80	19-APR-87
ADAMS	1100	23-MAY-87	19-APR-87
JONES	2975	02-APR-81	19-APR-87
FORD	3000	03-DEC-81	19-APR-87
SCOTT	3000	19-APR-87	19-APR-87

This example illustrates the nondeterministic nature of the LAST\_VALUE function. Ford and Scott have the same salary, so they are in adjacent rows. Ford appears first because the rows in the subquery are ordered by HIREDATE. However, if the rows are ordered by hiredate in descending order, as in the next example, the function returns a different value:

```

SELECT ename, sal, hiredate, LAST_VALUE(hiredate) OVER
  (ORDER BY sal
   ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING) AS lv
FROM (SELECT * FROM emp WHERE deptno=20 ORDER BY hiredate DESC);

```

ENAME	SAL	HIREDATE	LV
SMITH	800	17-DEC-80	19-APR-87
ADAMS	1100	23-MAY-87	19-APR-87
JONES	2975	02-APR-81	19-APR-87
FORD	3000	03-DEC-81	19-APR-87
SCOTT	3000	19-APR-87	19-APR-87

```

-----
SMITH          800 17-DEC-80 03-DEC-81
ADAMS         1100 23-MAY-87 03-DEC-81
JONES        2975 02-APR-81 03-DEC-81
SCOTT        3000 19-APR-87 03-DEC-81
FORD         3000 03-DEC-81 03-DEC-81

```

The following two examples show how to make the LAST\_VALUE function deterministic by ordering on a unique key. By ordering within the function by both salary and hiredate, you can ensure the same result regardless of the ordering in the subquery.

```

SELECT ename, sal, hiredate, LAST_VALUE(hiredate) OVER
       (ORDER BY sal, hiredate
        ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING) AS lv
FROM (SELECT * FROM emp WHERE deptno=20 ORDER BY hiredate);

```

```

ENAME          SAL HIREDATE  LV
-----
SMITH          800 17-DEC-80 19-APR-87
ADAMS         1100 23-MAY-87 19-APR-87
JONES        2975 02-APR-81 19-APR-87
FORD         3000 03-DEC-81 19-APR-87
SCOTT        3000 19-APR-87 19-APR-87

```

```

SELECT ename, sal, hiredate, LAST_VALUE(hiredate) OVER
       (ORDER BY sal, hiredate
        ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING) AS lv
FROM (SELECT * FROM emp WHERE deptno=20 ORDER BY hiredate DESC);

```


```

ENAME          SAL HIREDATE  LV
-----
SMITH          800 17-DEC-80 19-APR-87
ADAMS         1100 23-MAY-87 19-APR-87
JONES        2975 02-APR-81 19-APR-87
FORD         3000 03-DEC-81 19-APR-87
SCOTT        3000 19-APR-87 19-APR-87

```

# LEAD

## Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

## Purpose

LEAD is an analytic function. It provides access to more than one row of a table at the same time without a self-join. Given a series of rows returned from a query and a position of the cursor, LEAD provides access to a row at a given physical offset beyond that position.

If you do not specify *offset*, its default is 1. The optional *default* value is returned if the offset goes beyond the scope of the table. If you do not specify *default*, its default value is null.

## Example

The following example provides, for each employee in the EMP table, the hiredate of the employee hired just after:


```
SELECT ename, hiredate,
       LEAD(hiredate, 1) OVER (ORDER BY hiredate) AS next_hire_date
FROM emp;
```

ENAME	HIREDATE	NEXT_HIRE
SMITH	17-DEC-80	20-FEB-81
ALLEN	20-FEB-81	22-FEB-81
WARD	22-FEB-81	02-APR-81
JONES	02-APR-81	01-MAY-81
BLAKE	01-MAY-81	09-JUN-81
CLARK	09-JUN-81	08-SEP-81
TURNER	08-SEP-81	28-SEP-81
MARTIN	28-SEP-81	17-NOV-81
KING	17-NOV-81	03-DEC-81
JAMES	03-DEC-81	03-DEC-81
FORD	03-DEC-81	23-JAN-82
MILLER	23-JAN-82	19-APR-87

```
SCOTT      19-APR-87 23-MAY-87
ADAMS     23-MAY-87
```

## LEAST

### Syntax



### Purpose

LEAST returns the least of the list of *exprs*. All *exprs* after the first are implicitly converted to the datatype of the first *expr* before the comparison. Oracle compares the *exprs* using nonpadded comparison semantics. If the value returned by this function is character data, its datatype is always VARCHAR2.


### Example

```
SELECT LEAST('HARRY', 'HARRIOT', 'HAROLD') "LEAST"
       FROM DUAL;
```

```
LEAST
-----
HAROLD
```

## LENGTH

### Syntax



### Purpose

LENGTH returns the length of *char* in characters. If *char* has datatype CHAR, the length includes all trailing blanks. If *char* is null, this function returns null.

### Example

```
SELECT LENGTH('CANDIDE') "Length in characters"
```

```

FROM DUAL;

Length in characters
-----
7

```

## LENGTHB

### Syntax

```

→ [LENGTHB] → ( → char → ) →

```

### Purpose

LENGTHB returns the length of *char* in bytes. If *char* is null, this function returns null. For a single-byte database character set, LENGTHB is equivalent to LENGTH.

### Example

This example assumes a double-byte database character set.

```

SELECT LENGTHB ('CANDIDE') "Length in bytes"
       FROM DUAL;

Length in bytes
-----
14

```

## LN

### Syntax

```

→ [LN] → ( → n → ) →

```

### Purpose

LN returns the natural logarithm of *n*, where *n* is greater than 0.

### Example

```

SELECT LN(95) "Natural log of 95" FROM DUAL;

```



```
Natural log of 95
-----
          4.55387689
```

## LOG

### Syntax



### Purpose

LOG returns the logarithm, base *m*, of *n*. The base *m* can be any positive number other than 0 or 1 and *n* can be any positive number.

### Example

```
SELECT LOG(10,100) "Log base 10 of 100" FROM DUAL;
```

```
Log base 10 of 100
-----
                    2
```

## LOWER

### Syntax



### Purpose

LOWER returns *char*, with all letters lowercase. The return value has the same datatype as the argument *char* (CHAR or VARCHAR2).

### Example


```
SELECT LOWER('MR. SCOTT MCMILLAN') "Lowercase"
       FROM DUAL;
```

```
Lowercase
```

```
-----
mr. scott mcmillan
```

## LPAD

### Syntax



### Purpose

LPAD returns *char1*, left-padded to length *n* with the sequence of characters in *char2*; *char2* defaults to a single blank. If *char1* is longer than *n*, this function returns the portion of *char1* that fits in *n*.

The argument *n* is the total length of the return value as it is displayed on your terminal screen. In most character sets, this is also the number of characters in the return value. However, in some multibyte character sets, the display length of a character string can differ from the number of characters in the string.


### Example

```
SELECT LPAD('Page 1',15,'*.*') "LPAD example"
      FROM DUAL;
```

```
LPAD example
-----
*.*.*.*.*Page 1
```

## LTRIM

### Syntax



## Purpose

LTRIM removes characters from the left of *char*, with all the leftmost characters that appear in *set* removed; *set* defaults to a single blank. If *char* is a character literal, you must enclose it in single quotes. Oracle begins scanning *char* from its first character and removes all characters that appear in *set* until reaching a character not in *set* and then returns the result.


## Example

```
SELECT LTRIM('xyxXxyLAST WORD', 'xy') "LTRIM example"
      FROM DUAL;
```

```
LTRIM example
-----
XxyLAST WORD
```

# MAKE\_REF

## Syntax



## Purpose

MAKE\_REF creates a REF to a row of an object view or a row in an object table whose object identifier is primary key based.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals* for more information about object views.

## Example


```
CREATE TABLE emp (eno NUMBER, ename VARCHAR2(20),
                  salary NUMBER, PRIMARY KEY (eno, ename));
CREATE TYPE emp_type AS OBJECT
  (eno NUMBER, ename CHAR(20), salary NUMBER);
CREATE VIEW emp_view OF emp_type
  WITH OBJECT IDENTIFIER (eno, ename)
  AS SELECT * FROM emp;
SELECT MAKE_REF(emp_view, 1, 'jack') FROM DUAL;
```

```
MAKE_REF(EMP_VIEW,1,'JACK')
```

```
-----
000067030A0063420D06E06F3C00C1E03400400B40DCB10000001C26010001000200
29000000000000F0600810100140100002A0007000A8401FE0000001F02C102146A61
636B2020202020202020202020202020202020202020202000000000000000000000
00000000
```

## MAX

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

### Purpose

MAX returns maximum value of *expr*. You can use it as an aggregate or analytic function.

If you specify DISTINCT, you can specify only the *query\_partition\_clause* of the *analytic\_clause*. The *ORDER\_BY\_clause* and *windowing\_clause* are not allowed.

**See Also:** ["Aggregate Functions"](#) on page 4-6

### Aggregate Example

```
SELECT MAX(sal) "Maximum" FROM emp;
```

```

Maximum
-----
      5000
```

### Analytic Example

The following example calculates, for each employee, the highest salary of the employees reporting to the same manager as the employee.

```
SELECT mgr, ename, sal,
```

```

MAX(sal) OVER (PARTITION BY mgr) AS mgr_max
FROM emp;

```

MGR	ENAME	SAL	MGR_MAX
7566	SCOTT	3000	3000
7566	FORD	3000	3000
7698	ALLEN	1600	1600
7698	WARD	1250	1600
7698	JAMES	950	1600
7698	TURNER	1500	1600
7698	MARTIN	1250	1600
7782	MILLER	1300	1300
7788	ADAMS	1100	1100
7839	JONES	2975	2975
7839	CLARK	2450	2975
7839	BLAKE	2850	2975
7902	SMITH	800	800
	KING	5000	5000

If you enclose this query in the parent query with a predicate, you can determine the employee who makes the highest salary in each department:

```


SELECT mgr, ename, sal
FROM (SELECT mgr, ename, sal,
MAX(sal) OVER (PARTITION BY mgr) AS rmax_sal
FROM emp)
WHERE sal = rmax_sal;

```

MGR	ENAME	SAL
7566	SCOTT	3000
7566	FORD	3000
7698	ALLEN	1600
7782	MILLER	1300
7788	ADAMS	1100
7839	JONES	2975
7902	SMITH	800
	KING	5000

# MIN

## Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

## Purpose

MIN returns minimum value of *expr*. You can use it as an aggregate or analytic function.

If you specify *DISTINCT*, you can specify only the *query\_partition\_clause* of the *analytic\_clause*. The *ORDER\_BY\_clause* and *windowing\_clause* are not allowed.

**See Also:** ["Aggregate Functions"](#) on page 4-6

## Aggregate Example

```
SELECT MIN(hiredate) "Earliest" FROM emp;
```

```
Earliest
-----
17-DEC-80
```

## Analytic Example

The following example determines, for each employee, the employees who were hired on or before the same date as the employee. It then determines the subset of employees reporting to the same manager as the employee, and returns the lowest salary in that subset.


```
SELECT mgr, ename, hiredate, sal,
       MIN(sal) OVER(PARTITION BY mgr ORDER BY hiredate
                     RANGE UNBOUNDED PRECEDING) as p_cmin
FROM emp;
```

MGR	ENAME	HIREDATE	SAL	P_CMIN
7566	FORD	03-DEC-81	3000	3000

7566	SCOTT	19-APR-87	3000	3000
7698	ALLEN	20-FEB-81	1600	1600
7698	WARD	22-FEB-81	1250	1250
7698	TURNER	08-SEP-81	1500	1250
7698	MARTIN	28-SEP-81	1250	1250
7698	JAMES	03-DEC-81	950	950
7782	MILLER	23-JAN-82	1300	1300
7788	ADAMS	23-MAY-87	1100	1100
7839	JONES	02-APR-81	2975	2975
7839	BLAKE	01-MAY-81	2850	2850
7839	CLARK	09-JUN-81	2450	2450
7902	SMITH	17-DEC-80	800	800
	KING	17-NOV-81	5000	5000

## MOD

### Syntax



### Purpose

MOD returns remainder of  $m$  divided by  $n$ . Returns  $m$  if  $n$  is 0.

### Example

```
SELECT MOD(11,4) "Modulus" FROM DUAL;
```

```

Modulus
-----
          3
  
```

This function behaves differently from the classical mathematical modulus function when  $m$  is negative. The classical modulus can be expressed using the MOD function with this formula:

$$m - n * \text{FLOOR}(m/n)$$

The following statement illustrates the difference between the MOD function and the classical modulus:

```
SELECT m, n, MOD(m, n),
m - n * FLOOR(m/n) "Classical Modulus"
```

```
FROM test_mod_table;
```

M	N	MOD(M,N)	Classical Modulus
11	4	3	3
11	-4	3	-1
-11	4	-3	1
-11	-4	-3	-3

See Also: ["FLOOR"](#) on page 4-35.

## MONTHS\_BETWEEN

### Syntax

```
MONTHS_BETWEEN ( d1 , d2 )
```

### Purpose

MONTHS\_BETWEEN returns number of months between dates *d1* and *d2*. If *d1* is later than *d2*, result is positive; if earlier, negative. If *d1* and *d2* are either the same days of the month or both last days of months, the result is always an integer. Otherwise Oracle calculates the fractional portion of the result based on a 31-day month and considers the difference in time components of *d1* and *d2*.

### Example

```
SELECT MONTHS_BETWEEN
       ( TO_DATE( '02-02-1995' , 'MM-DD-YYYY' ) ,
         TO_DATE( '01-01-1995' , 'MM-DD-YYYY' ) ) "Months"
FROM DUAL;
```

```
Months
-----
1.03225806
```



## NEW\_TIME

### Syntax

→ NEW\_TIME ( ( d , z1 , z2 ) ) →

### Purpose

NEW\_TIME returns the date and time in time zone *z2* when date and time in time zone *z1* are *d*. The arguments *z1* and *z2* can be any of these text strings:

- AST, ADT: Atlantic Standard or Daylight Time
- BST, BDT: Bering Standard or Daylight Time
- CST, CDT: Central Standard or Daylight Time
- EST, EDT: Eastern Standard or Daylight Time
- GMT: Greenwich Mean Time
- HST, HDT: Alaska-Hawaii Standard Time or Daylight Time.
- MST, MDT: Mountain Standard or Daylight Time
- NST: Newfoundland Standard Time
- PST, PDT: Pacific Standard or Daylight Time
- YST, YDT: Yukon Standard or Daylight Time

## NEXT\_DAY

### Syntax

→ NEXT\_DAY ( ( d , char ) ) →

### Purpose

Returns the date of the first weekday named by *char* that is later than the date *d*. The argument *char* must be a day of the week in the date language of your session, either the full name or the abbreviation. The minimum number of letters required is the number of letters in the abbreviated version. Any characters immediately

following the valid abbreviation are ignored. The return value has the same hours, minutes, and seconds component as the argument *d*.

### Example

This example returns the date of the next Tuesday after March 15, 1998.

```
SELECT NEXT_DAY('15-MAR-98', 'TUESDAY') "NEXT DAY"
       FROM DUAL;

NEXT DAY
-----
16-MAR-98
```

## NLS\_CHARSET\_DECL\_LEN

### Syntax

```
→ NLS_CHARSET_DECL_LEN ( ( bytecnt , csid ) →
```

### Purpose

NLS\_CHARSET\_DECL\_LEN returns the declaration width (in number of characters) of an NCHAR column. The *bytecnt* argument is the width of the column. The *csid* argument is the character set ID of the column.

### Example

```
SELECT NLS_CHARSET_DECL_LEN
       (200, nls_charset_id('ja16eucfixed'))
       FROM DUAL;

NLS_CHARSET_DECL_LEN(200,NLS_CHARSET_ID('JA16EUCFIXED'))
-----
100
```

## NLS\_CHARSET\_ID

### Syntax

```
→ NLS_CHARSET_ID ( ( text ) →
```

## Purpose

NLS\_CHARSET\_ID returns the NLS character set ID number corresponding to NLS character set name, *text*. The *text* argument is a run-time VARCHAR2 value. The *text* value 'CHAR\_CS' returns the database character set ID number of the server. The *text* value 'NCHAR\_CS' returns the national character set ID number of the server.

Invalid character set names return null.

## Example

```
SELECT NLS_CHARSET_ID('ja16euc')
       FROM DUAL;
```

```
NLS_CHARSET_ID('JA16EUC')
-----
                        830
```

**See Also:** *Oracle8i National Language Support Guide* for a list of character set names.

# NLS\_CHARSET\_NAME

## Syntax

→ NLS\_CHARSET\_NAME ( ( n ) ) →

## Purpose

NLS\_CHARSET\_NAME returns the name of the NLS character set corresponding to ID number *n*. The character set name is returned as a VARCHAR2 value in the database character set.

If *n* is not recognized as a valid character set ID, this function returns null.

## Example


```
SELECT NLS_CHARSET_NAME ( 2 )
       FROM DUAL;
```

```
NLS_CH
-----
WE8DEC
```

**See Also:** *Oracle8i National Language Support Guide* for a list of character set IDs.

## NLS\_INITCAP

### Syntax



### Purpose

NLS\_INITCAP returns *char*, with the first letter of each word in uppercase, all other letters in lowercase. Words are delimited by white space or characters that are not alphanumeric. The value of '*nlsparam*' can have this form:

```
'NLS_SORT = sort'
```

where *sort* is either a linguistic sort sequence or BINARY. The linguistic sort sequence handles special linguistic requirements for case conversions. These requirements can result in a return value of a different length than the *char*. If you omit '*nlsparam*', this function uses the default sort sequence for your session.

### Example


```
SELECT NLS_INITCAP
       ('ijsland', 'NLS_SORT = XDutch') "Capitalized"
FROM DUAL;
```

```
Capital
-----
IJsland
```

**See Also:** *Oracle8i National Language Support Guide* for information on sort sequences.

## NLS\_LOWER

### Syntax



### Purpose

NLS\_LOWER returns *char*, with all letters lowercase. The '*nlsparm*' can have the same form and serve the same purpose as in the NLS\_INITCAP function.


### Example

```
SELECT NLS_LOWER
       ('CITTA''', 'NLS_SORT = XGerman') "Lowercase"
FROM DUAL;
```

```
Lower
-----
cittá
```

## NLSSORT

### Syntax



### Purpose

NLSSORT returns the string of bytes used to sort *char*. The value of '*nlsparms*' can have the form

```
'NLS_SORT = sort'
```

where *sort* is a linguistic sort sequence or BINARY. If you omit '*nlsparms*', this function uses the default sort sequence for your session. If you specify BINARY, this function returns *char*.

## Example

This function can be used to specify comparisons based on a linguistic sort sequence rather than on the binary value of a string:


```
SELECT ename FROM emp
   WHERE NLSSORT (ename, 'NLS_SORT = German')
   > NLSSORT ('S', 'NLS_SORT = German') ORDER BY ename;
```

```
ENAME
-----
SCOTT
SMITH
TURNER
WARD
```

**See Also:** *Oracle8i National Language Support Guide* for information on sort sequences.

## NLS\_UPPER

### Syntax



### Purpose

NLS\_UPPER returns *char*, with all letters uppercase. The '*nlsparm*' can have the same form and serve the same purpose as in the NLS\_INITCAP function.

### Example


```
SELECT NLS_UPPER
   ('große', 'NLS_SORT = XGerman') "Uppercase"
   FROM DUAL;
```

```
Upper
-----
GROSS
```

**See Also:** ["NLS\\_INITCAP"](#) on page 4-58.

# NTILE

## Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

## Purpose

NTILE is an analytic function. It divides an ordered dataset into a number of buckets indicated by *expr* and assigns the appropriate bucket number to each row. The buckets are numbered 1 through *expr*; and *expr* must resolve to a positive constant for each partition.

The number of rows in the buckets can differ by at most 1. The remainder values (the remainder of number of rows divided by buckets) are distributed 1 per bucket, starting with bucket 1.

If *expr* is greater than the number of rows, a number of buckets equal to the number of rows will be filled, and the remaining buckets will be empty.

## Example

The following example divides the values in the SAL column into 4 buckets. The SAL column has 14 values, so the two extra values (the remainder of 14 / 4) are allocated to buckets 1 and 2, which therefore have one more value than buckets 3 or 4.

```
SELECT ename, sal, NTILE(4) OVER (ORDER BY sal DESC) AS quartile
FROM emp;
```

ENAME	SAL	QUARTILE
KING	5000	1
SCOTT	3000	1
FORD	3000	1
JONES	2975	1
BLAKE	2850	2
CLARK	2450	2
ALLEN	1600	2
TURNER	1500	2

MILLER	1300	3
WARD	1250	3
MARTIN	1250	3
ADAMS	1100	4
JAMES	950	4
SMITH	800	4

## NUMTODSINTERVAL

---



---


**Note:** This function is restricted to use with analytic functions. It accepts only numbers as arguments, and returns interval literals. See "[Analytic Functions](#)" on page 4-7 and "[Interval](#)" on page 2-5.

---



---

### Syntax



### Purpose

NUMTODSINTERVAL converts *n* to an INTERVAL DAY TO SECOND literal. *n* can be a number or an expression resolving to a number. The value for *char\_expr* specifies the unit of *n* and must resolve to one of the following string values:

- 'DAY'
- 'HOUR'
- 'MINUTE'
- 'SECOND'

*char\_expr* is case-insensitive. Leading and trailing values within the parentheses are ignored. By default, precision of the return is 9.

### Example

The following example calculates for each employee, the number of employees hired, by the same manager, within the last 100 days from his/her hiredate:

```
SELECT mgr, ename, hiredate,
       COUNT(*) OVER (PARTITION BY mgr ORDER BY hiredate
                     RANGE NUMTODSINTERVAL(100, 'day') PRECEDING) AS t_count
FROM emp;
```



MGR	ENAME	HIREDATE	T_COUNT
7566	FORD	03-DEC-81	1
7566	SCOTT	19-APR-87	1
7698	ALLEN	20-FEB-81	1
7698	WARD	22-FEB-81	2
7698	TURNER	08-SEP-81	1
7698	MARTIN	28-SEP-81	2
7698	JAMES	03-DEC-81	3
7782	MILLER	23-JAN-82	1
7788	ADAMS	23-MAY-87	1
7839	JONES	02-APR-81	1
7839	BLAKE	01-MAY-81	2
7839	CLARK	09-JUN-81	3
7902	SMITH	17-DEC-80	1
	KING	17-NOV-81	1

## NUMTOYMINTERVAL

---

**Note:** This function is restricted to use with analytic functions. It accepts only numbers as arguments, and returns interval literals. See "[Analytic Functions](#)" on page 4-7 and "[Interval](#)" on page 2-5.

---

### Syntax

→ NUMTOYMINTERVAL ( ( n , ' char\_expr ' ) ) →

### Purpose

NUMTOYMINTERVAL converts number *n* to an INTERVAL YEAR TO MONTH literal. *n* can be a number or an expression resolving to a number. The value for *char\_expr* specifies the unit of *n*, and must resolve to one of the following string values:

- 'YEAR'
- 'MONTH'

*char\_expr* is case-insensitive. Leading and trailing values within the parentheses are ignored. By default, precision of the return is 9.

## Example


The following example calculates, for each employee, the total salary of employees hired in the past one year from his/her hiredate.

```
SELECT ename, hiredate, sal, SUM(sal) OVER (ORDER BY hiredate
      RANGE NUMTOYMINTERVAL(1,'year') PRECEDING) AS t_sal
FROM emp;
```

ENAME	HIREDATE	SAL	T_SAL
SMITH	17-DEC-80	800	800
ALLEN	20-FEB-81	1600	2400
WARD	22-FEB-81	1250	3650
JONES	02-APR-81	2975	6625
BLAKE	01-MAY-81	2850	9475
CLARK	09-JUN-81	2450	11925
TURNER	08-SEP-81	1500	13425
MARTIN	28-SEP-81	1250	14675
KING	17-NOV-81	5000	19675
JAMES	03-DEC-81	950	23625
FORD	03-DEC-81	3000	23625
MILLER	23-JAN-82	1300	24125
SCOTT	19-APR-87	3000	3000
ADAMS	23-MAY-87	1100	4100

## NVL

### Syntax



### Purpose

If *expr1* is null, NVL returns *expr2*; if *expr1* is not null, NVL returns *expr1*. The arguments *expr1* and *expr2* can have any datatype. If their datatypes are different, Oracle converts *expr2* to the datatype of *expr1* before comparing them. The datatype of the return value is always the same as the datatype of *expr1*, unless *expr1* is character data, in which case the return value's datatype is VARCHAR2.

### Example


```
SELECT ename, NVL(TO_CHAR(COMM), 'NOT APPLICABLE')
```

```
"COMMISSION" FROM emp
WHERE deptno = 30;
```

ENAME	COMMISSION
ALLEN	300
WARD	500
MARTIN	1400
BLAKE	NOT APPLICABLE
TURNER	0
JAMES	NOT APPLICABLE

## NVL2

### Syntax



### Purpose

If *expr1* is not null, NVL2 returns *expr2*. If *expr1* is null, NVL2 returns *expr3*. The argument *expr1* can have any datatype. The arguments *expr2* and *expr3* can have any datatypes except LONG.

If the datatypes of *expr2* and *expr3* are different, Oracle converts *expr3* to the datatype of *expr2* before comparing them unless *expr3* is a null constant. In that case, a datatype conversion is not necessary.

The datatype of the return value is always the same as the datatype of *expr2*, unless *expr2* is character data, in which case the return value's datatype is VARCHAR2.

### Example

The following example shows whether the income of each employee in department 30 is made up of salary plus commission, or just salary, depending on whether the COMM column of EMP is null or not.


```
SELECT ename, NVL2(TO_CHAR(COMM), 'SAL & COMM', 'SAL') income
FROM emp WHERE deptno = 30;
```

ENAME	INCOME
ALLEN	SAL & COMM

WARD	SAL & COMM
MARTIN	SAL & COMM
BLAKE	SAL
TURNER	SAL & COMM
JAMES	SAL

## PERCENT\_RANK

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

PERCENT\_RANK is an analytic function, and is similar to the CUME\_DIST (cumulative distribution) function. For a row R, PERCENT\_RANK calculates the rank of R minus 1, divided by 1 less than the number of rows being evaluated (the entire query result set or a partition). The range of values returned by PERCENT\_RANK is 0 to 1, inclusive. The first row in any set has a PERCENT\_RANK of 0.

### Example

The following example calculates, for each employee, the percent rank of the employee's salary within the department:

```

SELECT deptno, ename, sal,
       PERCENT_RANK() OVER (PARTITION BY deptno ORDER BY sal DESC) AS pr
FROM emp;


```

DEPTNO	ENAME	SAL	PR
10	KING	5000	0
10	CLARK	2450	.5
10	MILLER	1300	1
20	SCOTT	3000	0
20	FORD	3000	0
20	JONES	2975	.5
20	ADAMS	1100	.75
20	SMITH	800	1

30	BLAKE	2850	0
30	ALLEN	1600	.2
30	TURNER	1500	.4
30	WARD	1250	.6
30	MARTIN	1250	.6
30	JAMES	950	1

## POWER

### Syntax



### Purpose

POWER returns  $m$  raised to the  $n$ th power. The base  $m$  and the exponent  $n$  can be any numbers, but if  $m$  is negative,  $n$  must be an integer.

### Example

```
SELECT POWER(3,2) "Raised" FROM DUAL;
```


```

Raised
-----
          9

```

## RANK

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

RANK is an analytic function. It computes the rank of each row returned from a query with respect to the other rows returned by the query, based on the values of the *value\_exprs* in the *ORDER\_BY\_clause*. Rows with equal values for the ranking

criteria receive the same rank. Oracle then adds the number of tied rows to the tied rank to calculate the next rank. Therefore, the ranks may not be consecutive numbers.

### Example


The following statement ranks the employees within each department based on their salary and commission. Identical salary values receive the same rank and cause nonconsecutive ranks. Compare this example with the example for ["DENSE\\_RANK"](#) on page 4-29.

```
SELECT deptno, ename, sal, comm,
       RANK() OVER (PARTITION BY deptno ORDER BY sal DESC, comm) as rk
FROM emp;
```

DEPTNO	ENAME	SAL	COMM	RK
10	KING	5000		1
10	CLARK	2450		2
10	MILLER	1300		3
20	SCOTT	3000		1
20	FORD	3000		1
20	JONES	2975		3
20	ADAMS	1100		4
20	SMITH	800		5
30	BLAKE	2850		1
30	ALLEN	1600	300	2
30	TURNER	1500	0	3
30	WARD	1250	500	4
30	MARTIN	1250	1400	5
30	JAMES	950		6

## RATIO\_TO\_REPORT

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

## Purpose

RATIO\_TO\_REPORT is an analytic function. It computes the ratio of a value to the sum of a set of values. If *expr* evaluates to null, the ratio-to-report value also evaluates to null.

The set of values is determined by the *query\_partition\_clause*. If you omit that clause, the ratio-to-report is computed over all rows returned by the query.

## Example

The following example calculates the ratio-to-report of each salesperson's salary to the total of all salespeople's salaries:

```
SELECT ename, sal, RATIO_TO_REPORT(sal) OVER () AS rr
       FROM emp
       WHERE job = 'SALESMAN';
```

ENAME	SAL	RR
ALLEN	1600	.285714286
WARD	1250	.223214286
MARTIN	1250	.223214286
TURNER	1500	.267857143

# RAWTOHEX

## Syntax

```
→ RAWTOHEX → ( → raw → ) →
```

## Purpose

RAWTOHEX converts *raw* to a character value containing its hexadecimal equivalent.


## Example

```
SELECT RAWTOHEX(raw_column) "Graphics"
       FROM graphics;
```

```
Graphics
-----
7D
```

## REF

### Syntax



The diagram shows the syntax for the REF function: a box containing 'REF' followed by a left parenthesis '(', a box containing 'correlation\_variable', and a right parenthesis ')'. Arrows indicate the flow from left to right between these elements.

### Purpose

In a SQL statement, REF takes as its argument a correlation variable (table alias) associated with a row of an object table or an object view. A REF value is returned for the object instance that is bound to the variable or row.

### Example

```
CREATE TYPE emp_type AS OBJECT
  (eno NUMBER, ename VARCHAR2(20), salary NUMBER);
CREATE TABLE emp_table OF emp_type
  (primary key (eno, ename));
INSERT INTO emp_table VALUES (10, 'jack', 50000);
SELECT REF(e) FROM emp_table e;
```


```
REF(E)
```

```
-----
0000280209420D2FEABD9400C3E03400400B40DCB1420D2FEABD9300C3E03400400B
40DCB1004049EE0000
```

**See Also:** *Oracle8i Concepts*.

## REFTOHEX

### Syntax



The diagram shows the syntax for the REFTOHEX function: a box containing 'REFTOHEX' followed by a left parenthesis '(', a box containing 'r', and a right parenthesis ')'. Arrows indicate the flow from left to right between these elements.

### Purpose

REFTOHEX converts argument *r* to a character value containing its hexadecimal equivalent.



**Example**

```

CREATE TYPE emp_type AS OBJECT
  (eno NUMBER, ename VARCHAR2(20), salary NUMBER);
CREATE TABLE emp_table OF emp_type
  (primary key (eno, ename));
CREATE TABLE dept
  (dno NUMBER, mgr REF emp_type SCOPE IS emp);
INSERT INTO emp_table VALUES (10, 'jack', 50000);
INSERT INTO dept SELECT 10, REF(e) FROM emp_table e;
SELECT REFTOHEX(mgr) FROM dept;

REFTOHEX(MGR)
-----
0000220208420D2FEABD9400C3E03400400B40DCB1420D2FEABD9300C3E03400400B
40DCB1

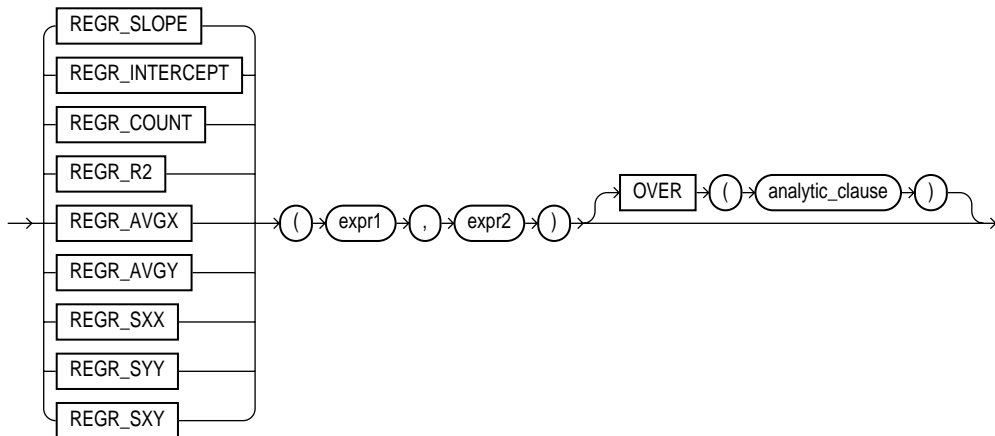
```

**REGR\_ (linear regression) functions**

The linear regression functions are:

- REGR\_SLOPE
- REGR\_INTERCEPT
- REGR\_COUNT
- REGR\_R2
- REGR\_AVGX
- REGR\_AVGY
- REGR\_SXX
- REGR\_SYY
- REGR\_SXY

## Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

## Purpose

The linear regression functions fit an ordinary-least-squares regression line to a set of number pairs. You can use them as both aggregate and analytic functions.

**See Also:** ["Aggregate Functions"](#) on page 4-6

Oracle applies the function to the set of (*expr1*, *expr2*) pairs after eliminating all pairs for which either *expr1* or *expr2* is null. Oracle computes all the regression functions simultaneously during a single pass through the data.

*expr1* is interpreted as a value of the dependent variable (a "y value"), and *expr2* is interpreted as a value of the independent variable (an "x value"). Both expressions must be numbers.

- REGR\_SLOPE returns the slope of the line. The return value is a number and can be null. After the elimination of null (*expr1*, *expr2*) pairs, it makes the following computation:

$$\text{COVAR\_POP}(\text{expr1}, \text{expr2}) / \text{VAR\_POP}(\text{expr2})$$

- REGR\_INTERCEPT returns the y-intercept of the regression line. The return value is a number and can be null. After the elimination of null (*expr1*, *expr2*) pairs, it makes the following computation:

$$\text{AVG}(\text{expr1}) - \text{REGR\_SLOPE}(\text{expr1}, \text{expr2}) * \text{AVG}(\text{expr2})$$

- REGR\_COUNT returns an integer that is the number of non-null number pairs used to fit the regression line.
- REGR\_R2 returns the coefficient of determination (also called "R-squared" or "goodness of fit") for the regression. The return value is a number and can be null. VAR\_POP(*expr1*) and VAR\_POP(*expr2*) are evaluated after the elimination of null pairs. The return values are:

$$\text{NULL if VAR\_POP}(\text{expr2}) = 0$$

$$1 \text{ if VAR\_POP}(\text{expr1}) = 0 \text{ and} \\ \text{VAR\_POP}(\text{expr2}) \neq 0$$

$$\text{POWER}(\text{CORR}(\text{expr1}, \text{expr2}), 2) \text{ if VAR\_POP}(\text{expr1}) > 0 \text{ and} \\ \text{VAR\_POP}(\text{expr2}) \neq 0$$

All of the remaining regression functions return a number and can be null:

- REGR\_AVGX evaluates the average of the independent variable (*expr2*) of the regression line. It makes the following computation after the elimination of null (*expr1*, *expr2*) pairs:

$$\text{AVG}(\text{expr2})$$

- REGR\_AVGY evaluates the average of the dependent variable (*expr1*) of the regression line. It makes the following computation after the elimination of null (*expr1*, *expr2*) pairs:

$$\text{AVG}(\text{expr1})$$

REGR\_SXY, REGR\_SXX, REGR\_SYY are auxiliary functions that are used to compute various diagnostic statistics.

- REGR\_SXX makes the following computation after the elimination of null (*expr1*, *expr2*) pairs:

$$\text{REGR\_COUNT}(\text{expr1}, \text{expr2}) * \text{VAR\_POP}(\text{expr2})$$

- REGR\_SYY makes the following computation after the elimination of null (*expr1*, *expr2*) pairs:

$$\text{REGR\_COUNT}(\text{expr1}, \text{expr2}) * \text{VAR\_POP}(\text{expr1})$$

- REGR\_SXY makes the following computation after the elimination of null (*expr1*, *expr2*) pairs:

```
REGR_COUNT(expr1, expr2) * COVAR_POP(expr1, expr2)
```

The following examples are based on the SALES table, described in "COVAR\_POP" on page 4-26.

### REGR\_SLOPE and REGR\_INTERCEPT Examples

The following example determines the slope and intercept of the regression line for the amount of sales and sale profits for each year.

```
SELECT s_year,
       REGR_SLOPE(s_amount, s_profit),
       REGR_INTERCEPT(s_amount, s_profit)
FROM sales GROUP BY s_year;
```

S_YEAR	REGR_SLOPE	REGR_INTER
1998	128.401558	-2277.5684
1999	55.618655	226.855296

The following example determines the cumulative slope and cumulative intercept of the regression line for the amount of sales and sale profits for each day in 1998:

```
SELECT s_year, s_month, s_day,
       REGR_SLOPE(s_amount, s_profit)
         OVER (ORDER BY s_month, s_day) AS CUM_SLOPE,
       REGR_INTERCEPT(s_amount, s_profit)
         OVER (ORDER BY s_month, s_day) AS CUM_ICPT
FROM sales
WHERE s_year=1998
ORDER BY s_month, s_day;
```

S_YEAR	S_MONTH	S_DAY	CUM_SLOPE	CUM_ICPT
1998		6	5	
1998		6	9	132.093066 401.884833
1998		6	9	132.093066 401.884833
1998		6	10	131.829612 450.65349
1998		8	21	132.963737 -153.5413
1998		8	25	130.681718 -451.47349
1998		8	25	130.681718 -451.47349
1998		8	26	128.76502 -236.50096
1998		11	9	131.499934 -1806.7535
1998		11	9	131.499934 -1806.7535
1998		11	10	130.190972 -2323.3056

1998	11	10	130.190972	-2323.3056
1998	11	11	128.401558	-2277.5684

### REGR\_COUNT Examples

The following example returns the number of sales transactions in the SALES table that resulted in a profit. (None of the rows for containing a sales amount have a null in the S\_PROFIT column, so the function returns the total number of rows in the SALES table.)

```
SELECT REGR_COUNT(s_amount, s_profit) FROM sales;

REGR_COUNT
-----
          23
```

The following example computes, for each day, the cumulative number of transactions within each month for the year 1998:

```
SELECT s_month, s_day,
       REGR_COUNT(s_amount,s_profit)
       OVER (PARTITION BY s_month ORDER BY s_day)
FROM SALES
WHERE S_YEAR=1998
ORDER BY S_MONTH;
```

S_MONTH	S_DAY	REGR_COUNT
6	5	1
6	9	3
6	9	3
6	10	4
8	21	1
8	25	3
8	25	3
8	26	4
11	9	2
11	9	2
11	10	4
11	10	4
11	11	5

### REGR\_R2 Examples

The following example computes the coefficient of determination of the regression line for amount of sales and sale profits:

```
SELECT REGR_R2(s_amount, s_profit) FROM sales;
```

```
REGR_R2(S_
-----
.942435028
```

The following example computes the cumulative coefficient of determination of the regression line for monthly sales and monthly profits for each month in 1998:

```
SELECT s_month,
       REGR_R2(SUM(s_amount), SUM(s_profit))
         OVER (ORDER BY s_month)
FROM SALES
WHERE s_year=1998
GROUP BY s_month
ORDER BY s_month;
```

```
S_MONTH      REGR_R2(SU
-----
           6
           8          1
          11 .740553632
```

### REGR\_AVGY and REGR\_AVGX Examples

The following example calculates the regression average for the amount of sales and sale profits for each year:

```
SELECT s_year,
       REGR_AVGY(s_amount, s_profit),
       REGR_AVGX(s_amount, s_profit)
FROM sales GROUP BY s_year;
```

```
S_YEAR      REGR_AVGY( REGR_AVGX(
-----
          1998 41227.5462 338.820769
          1999  7330.748   127.725
```

The following example calculates the cumulative averages for the amount of sales and sale profits in 1998:

```
SELECT s_year, s_month, s_day,
       REGR_AVGY(s_amount, s_profit)
         OVER (ORDER BY s_month, s_day) AS CUM_AMOUNT,
       REGR_AVGX(s_amount, s_profit)
         OVER (ORDER BY s_month, s_day) AS CUM_PROFIT
```

```
FROM sales
WHERE s_year=1998
ORDER BY s_month, s_day;
```

S_YEAR	S_MONTH	S_DAY	CUM_AMOUNT	CUM_PROFIT
1998	6	5	16068	118.2
1998	6	9	44375.6667	332.9
1998	6	9	44375.6667	332.9
1998	6	10	52678.25	396.175
1998	8	21	44721.72	337.5
1998	8	25	45333.8	350.357143
1998	8	25	45333.8	350.357143
1998	8	26	47430.7	370.1875
1998	11	9	41892.91	332.317
1998	11	9	41892.91	332.317
1998	11	10	40777.175	331.055833
1998	11	10	40777.175	331.055833
1998	11	11	41227.5462	338.820769

### REGR\_SXY, REGR\_SXX, and REGR\_SYY Examples

The following example computes the REGR\_SXY, REGR\_SXX, and REGR\_SYY values for the regression analysis of amount of sales and sale profits for each year:

```
SELECT s_year,
       REGR_SXY(s_amount, s_profit),
       REGR_SYY(s_amount, s_profit),
       REGR_SXX(s_amount, s_profit)
FROM sales GROUP BY s_year;
```

S_YEAR	REGR_SXY(S	REGR_SYY(S	REGR_SXX(S
1998	48723551.8	6423698688	379462.311
1999	3605361.62	200525751	64822.8841

The following example computes the cumulative REGR\_SXY, REGR\_SXX, and REGR\_SYY statistics for amount of sales and sale profits for each month-day value in 1998:

```
SELECT s_year, s_month, s_day,
       REGR_SXY(s_amount, s_profit)
       OVER (ORDER BY s_month, s_day) AS CUM_SXY,
       REGR_SYY(s_amount, s_profit)
       OVER (ORDER BY s_month, s_day) AS CUM_SYY,
       REGR_SXX(s_amount, s_profit)
```

## REPLACE


---

```
                OVER (ORDER BY s_month, s_day) AS CUM_SXX
FROM sales
WHERE s_year=1998
ORDER BY s_month, s_day;
```

S_YEAR	S_MONTH	S_DAY	CUM_SXY	CUM_SXY	CUM_SXX
1998	6	5	0	0	0
1998	6	9	14822857.8	1958007601	112215.26
1998	6	9	14822857.8	1958007601	112215.26
1998	6	10	21127009.3	2785202281	160259.968
1998	8	21	30463997.3	4051329674	229115.08
1998	8	25	34567985.3	4541739739	264520.437
1998	8	25	34567985.3	4541739739	264520.437
1998	8	26	36896592.7	4787971157	286542.049
1998	11	9	45567995.3	6045196901	346524.854
1998	11	9	45567995.3	6045196901	346524.854
1998	11	10	48178003.8	6392056557	370056.411
1998	11	10	48178003.8	6392056557	370056.411
1998	11	11	48723551.8	6423698688	379462.311

## REPLACE

### Syntax



### Purpose

REPLACE returns *char* with every occurrence of *search\_string* replaced with *replacement\_string*. If *replacement\_string* is omitted or null, all occurrences of *search\_string* are removed. If *search\_string* is null, *char* is returned. This function provides a superset of the functionality provided by the TRANSLATE function. TRANSLATE provides single-character, one-to-one substitution. REPLACE lets you substitute one string for another as well as to remove character strings.

### Example

```
SELECT REPLACE('JACK and JUE', 'J', 'BL') "Changes"
FROM DUAL;
```




Changes

-----  
BLACK and BLUE

## ROUND (Number Function)

### Syntax



### Purpose

ROUND returns  $n$  rounded to  $m$  places right of the decimal point. If  $m$  is omitted,  $n$  is rounded to 0 places.  $m$  can be negative to round off digits left of the decimal point.  $m$  must be an integer.

### Example 1

```
SELECT ROUND(15.193,1) "Round" FROM DUAL;
```

```
Round
-----
    15.2
```


### Example 2

```
SELECT ROUND(15.193,-1) "Round" FROM DUAL;
```

```
Round
-----
    20
```

## ROUND (Date Function)

### Syntax



## Purpose

ROUND returns *d* rounded to the unit specified by the format model *fmt*. If you omit *fmt*, *d* is rounded to the nearest day. See ["ROUND and TRUNC Date Functions"](#) on page 4-117 for the permitted format models to use in *fmt*.


## Example

```
SELECT ROUND (TO_DATE ('27-OCT-92'), 'YEAR')
         "New Year" FROM DUAL;
```

```
New Year
-----
01-JAN-93
```

# ROW\_NUMBER

## Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

## Purpose

ROW\_NUMBER is an analytic function. It assigns a unique number to each row to which it is applied (either each row in the partition or each row returned by the query), in the ordered sequence of rows specified in the *ORDER\_BY\_clause*, beginning with 1.

## Example

For each department in the EMP table, the following example assigns numbers to each row in order of employee's hire date:

```
SELECT deptno, ename, hiredate, ROW_NUMBER()
       OVER (PARTITION BY deptno ORDER BY hiredate) AS emp_id
       FROM emp;
```

```
DEPTNO ENAME      HIREDATE      EMP_ID
-----
      10 CLARK      09-JUN-81      1
```

10	KING	17-NOV-81	2
10	MILLER	23-JAN-82	3
20	SMITH	17-DEC-80	1
20	JONES	02-APR-81	2
20	FORD	03-DEC-81	3
20	SCOTT	19-APR-87	4
20	ADAMS	23-MAY-87	5
30	ALLEN	20-FEB-81	1
30	WARD	22-FEB-81	2
30	BLAKE	01-MAY-81	3
30	TURNER	08-SEP-81	4
30	MARTIN	28-SEP-81	5
30	JAMES	03-DEC-81	6

ROW\_NUMBER is a nondeterministic function. However, HIREDATE is a unique key, so the results of this application of the function are deterministic. For examples of nondeterministic behavior, see "FIRST\_VALUE" on page 4-34 and "LAST\_VALUE" on page 4-42.

## ROWIDTOCHAR

### Syntax

→ ROWIDTOCHAR ( ( rowid ) ) →

### Purpose

ROWIDTOCHAR converts a rowid value to VARCHAR2 datatype. The result of this conversion is always 18 characters long.


### Example

```
SELECT ROWID
   FROM offices
  WHERE
     ROWIDTOCHAR(ROWID) LIKE '%Br1AAB%';
```

```
ROWID
-----
AAAAZ6AABAAABr1AAB
```

# RPAD

## Syntax



## Purpose

RPAD returns *char1*, right-padded to length *n* with *char2*, replicated as many times as necessary; *char2* defaults to a single blank. If *char1* is longer than *n*, this function returns the portion of *char1* that fits in *n*.

The argument *n* is the total length of the return value as it is displayed on your terminal screen. In most character sets, this is also the number of characters in the return value. However, in some multibyte character sets, the display length of a character string can differ from the number of characters in the string.


## Example

```
SELECT RPAD('MORRISON',12,'ab') "RPAD example"
      FROM DUAL;
```

```
RPAD example
-----
MORRISONabab
```

# RTRIM

## Syntax



## Purpose

RTRIM returns *char*, with all the rightmost characters that appear in *set* removed; *set* defaults to a single blank. If *char* is a character literal, you must enclose it in single quotes. RTRIM works similarly to LTRIM.

**Example**

```
SELECT RTRIM('BROWNINGyxXxy', 'xy') "RTRIM e.g."
       FROM DUAL;
```

```
RTRIM e.g
-----
BROWNINGyxX
```

**See Also:** ["LTRIM"](#) on page 4-48.

## SIGN

**Syntax**

```
→ SIGN ( n ) →
```

**Purpose**

If  $n < 0$ , SIGN returns -1. If  $n = 0$ , the function returns 0. If  $n > 0$ , SIGN returns 1.

**Example**

```
SELECT SIGN(-15) "Sign" FROM DUAL;
```

```
      Sign
-----
      -1
```

## SIN

**Syntax**

```
→ SIN ( n ) →
```

**Purpose**

SIN returns the sine of  $n$  (an angle expressed in radians).

**Example**

```
SELECT SIN(30 * 3.14159265359/180)
       "Sine of 30 degrees" FROM DUAL;
```

```
Sine of 30 degrees
-----
                        .5
```

**SINH****Syntax**

→ **SINH** → ( → *n* → ) →

**Purpose**

SINH returns the hyperbolic sine of *n*.

**Example**

```
SELECT SINH(1) "Hyperbolic sine of 1" FROM DUAL;
```

```
Hyperbolic sine of 1
-----
                1.17520119
```

**SOUNDEX****Syntax**

→ **SOUNDEX** → ( → *char* → ) →

**Purpose**

SOUNDEX returns a character string containing the phonetic representation of *char*. This function allows you to compare words that are spelled differently, but sound alike in English.

The phonetic representation is defined in *The Art of Computer Programming, Volume 3: Sorting and Searching*, by Donald E. Knuth, as follows:

- Retain the first letter of the string and remove all other occurrences of the following letters: a, e, h, i, o, u, w, y.
- Assign numbers to the remaining letters (after the first) as follows:
  - b, f, p, v = 1
  - c, g, j, k, q, s, x, z = 2
  - d, t = 3
  - l = 4
  - m, n = 5
  - r = 6
- If two or more letters with the same number were adjacent in the original name (before step 1), or adjacent except for any intervening h and w, omit all but the first.
- Return the first four bytes padded with 0.

### Example

```
SELECT ename
       FROM emp
      WHERE SOUNDEX(ename)
            = SOUNDEX('SMYTHE');
```

```
ENAME
-----
SMITH
```

## SQRT

### Syntax

→ **SQRT** ( *n* ) →

### Purpose

SQRT returns square root of *n*. The value *n* cannot be negative. SQRT returns a "real" result.

### Example

```
SELECT SQRT(26) "Square root" FROM DUAL;
```


```

Square root
-----
5.09901951

```

## STDDEV

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

### Purpose

STDDEV returns sample standard deviation of *expr*, a set of numbers. You can use it as both an aggregate and analytic function. It differs from STDDEV\_SAMP in that STDDEV returns zero when it has only 1 row of input data, whereas STDDEV\_SAMP returns a null.

Oracle calculates the standard deviation as the square root of the variance defined for the VARIANCE aggregate function.

If you specify DISTINCT, you can specify only the *query\_partition\_clause* of the *analytic\_clause*. The *ORDER\_BY\_clause* and *windowing\_clause* are not allowed.

**See Also:** ["Aggregate Functions"](#) on page 4-6, ["VARIANCE"](#) on page 4-115, and ["STDDEV\\_SAMP"](#) on page 4-88.

### Aggregate Example

```

SELECT STDDEV(sal) "Deviation"
   FROM emp;

```

```

Deviation
-----
1182.50322

```



## Analytic Example


The query in the following example returns the cumulative standard deviation of salary values in Department 30 ordered by hiredate:

```
SELECT ENAME, SAL, STDDEV(SAL) OVER (ORDER BY HIREDATE)
FROM EMP
WHERE DEPTNO=30;
```

ENAME	SAL	STDDEV(SAL)
ALLEN	1600	0
WARD	1250	247.487373
BLAKE	2850	841.130192
TURNER	1500	715.308791
MARTIN	1250	666.520817
JAMES	950	668.331255

## STDDEV\_POP

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

### Purpose

STDDEV\_POP computes the population standard deviation and returns the square root of the population variance. You can use it as both an aggregate and analytic function.

The *expr* is a number expression, and the function returns a value of type NUMBER. This function is same as the square root of the VAR\_POP function. When VAR\_POP returns null, this function returns null.

**See Also:** ["Aggregate Functions"](#) on page 4-6 and ["VAR\\_POP"](#) on page 4-113.

### Aggregate Example

The following example returns the population and sample standard deviations of profit from sales in the SALES table.

```
SELECT STDDEV_POP(s_profit), STDDEV_SAMP(s_profit) FROM sales;
```

```
STDDEV_POP STDDEV_SAMP
-----
173.975774 177.885831
```

### Analytic Example


The following example returns the population standard deviations of salaries in the EMP table by department:

```
SELECT deptno, ename, sal,
       STDDEV_POP(sal) OVER (PARTITION BY deptno) AS pop_std
FROM emp;
```

DEPTNO	ENAME	SAL	POP_STD
10	CLARK	2450	1546.14215
10	KING	5000	1546.14215
10	MILLER	1300	1546.14215
20	SMITH	800	1004.73877
20	ADAMS	1100	1004.73877
20	FORD	3000	1004.73877
20	SCOTT	3000	1004.73877
20	JONES	2975	1004.73877
30	ALLEN	1600	610.100174
30	BLAKE	2850	610.100174
30	MARTIN	1250	610.100174
30	JAMES	950	610.100174
30	TURNER	1500	610.100174
30	WARD	1250	610.100174

## STDDEV\_SAMP

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

STDDEV\_SAMP computes the cumulative sample standard deviation and returns the square root of the sample variance. You can use it as both an aggregate and analytic function.

The *expr* is a number expression, and the function returns a value of type NUMBER. This function is same as the square root of the VAR\_SAMP function. When VAR\_SAMP returns null, this function returns null.

**See Also:** "[Aggregate Functions](#)" on page 4-6 and "[VAR\\_SAMP](#)" on page 4-114.

### Aggregate Example

See the example for "[STDDEV\\_POP](#)" on page 4-87.

### Analytic Example


The following example returns the sample standard deviation of salaries in the EMP table by department:

```
SELECT deptno, ename, hiredate, sal,
       STDDEV_SAMP(sal) OVER (PARTITION BY deptno ORDER BY hiredate
                              ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW) AS cum_sdev
FROM emp;
```

DEPTNO	ENAME	HIREDATE	SAL	CUM_SDEV
10	CLARK	09-JUN-81	2450	
10	KING	17-NOV-81	5000	1803.12229
10	MILLER	23-JAN-82	1300	1893.62967
20	SMITH	17-DEC-80	800	
20	JONES	02-APR-81	2975	1537.95725
20	FORD	03-DEC-81	3000	1263.01557
20	SCOTT	19-APR-87	3000	1095.8967
20	ADAMS	23-MAY-87	1100	1123.3321
30	ALLEN	20-FEB-81	1600	
30	WARD	22-FEB-81	1250	247.487373
30	BLAKE	01-MAY-81	2850	841.130192
30	TURNER	08-SEP-81	1500	715.308791
30	MARTIN	28-SEP-81	1250	666.520817
30	JAMES	03-DEC-81	950	668.331255

# SUBSTR

## Syntax



## Purpose

SUBSTR returns a portion of *char*, beginning at character *m*, *n* characters long.

- If *m* is 0, it is treated as 1.
- If *m* is positive, Oracle counts from the beginning of *char* to find the first character.
- If *m* is negative, Oracle counts backwards from the end of *char*.
- If *n* is omitted, Oracle returns all characters to the end of *char*. If *n* is less than 1, a null is returned.

Floating-point numbers passed as arguments to SUBSTR are automatically converted to integers.

```
SELECT SUBSTR('ABCDEFGF',3,4) "Substring"
       FROM DUAL;
```

```
Substring
-----
CDEF
```


## Example 2

```
SELECT SUBSTR('ABCDEFGF',-5,4) "Substring"
       FROM DUAL;
```

```
Substring
-----
CDEF
```

## SUBSTRB

### Syntax



### Purpose

SUBSTRB is the same as SUBSTR, except that the arguments *m* and *n* are expressed in bytes, rather than in characters. For a single-byte database character set, SUBSTRB is equivalent to SUBSTR.

Floating-point numbers passed as arguments to SUBSTRB are automatically converted to integers.

### Example


Assume a double-byte database character set:

```
SELECT SUBSTRB('ABCDEFGF',5,4.2)
       "Substring with bytes"
FROM DUAL;
```

```
Substring with bytes
-----
CD
```

## SUM

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

## Purpose

SUM returns sum of values of *expr*. You can use it as an aggregate or analytic function.

**See Also:** ["Aggregate Functions"](#) on page 4-6

## Aggregate Function Example

The following example calculates the sum of all salaries in the EMP table:

```
SELECT SUM(sal) "Total"
       FROM emp;
```

```
       Total
-----
       29025
```

## Analytic Example


The following example calculates, for each manager, a cumulative total of salaries of employees who answer to that manager that are equal to or less than the current salary:

```
SELECT mgr, ename, sal,
       SUM(sal) OVER (PARTITION BY mgr ORDER BY sal
                     RANGE UNBOUNDED PRECEDING) l_csum
       FROM emp;
```

MGR	ENAME	SAL	L_CSUM
7566	SCOTT	3000	6000
7566	FORD	3000	6000
7698	JAMES	950	950
7698	WARD	1250	3450
7698	MARTIN	1250	3450
7698	TURNER	1500	4950
7698	ALLEN	1600	6550
7782	MILLER	1300	1300
7788	ADAMS	1100	1100
7839	CLARK	2450	2450
7839	BLAKE	2850	5300
7839	JONES	2975	8275
7902	SMITH	800	800
	KING	5000	5000

## SYS\_CONTEXT

### Syntax



### Purpose

`SYS_CONTEXT` returns the value of *attribute* associated with the context *namespace*. You can use this function in both SQL and PL/SQL statements. The context *namespace* must already have been created, and the associated *attribute* and its value must also have been set using the `DBMS_SESSION.SET_CONTEXT` procedure. The *namespace* must be a valid SQL identifier, and the *attribute* name cannot exceed 30 bytes in length.

The datatype of the return value is `VARCHAR2`. The default maximum size of the return value is 256 bytes. You can override this default by specifying the optional *length* parameter. The valid range of values is 1 to 4000 bytes. (If you specify an invalid value, Oracle ignores it and uses the default.)

Oracle8i provides a built-in namespace called `USERENV`, which describes the current session. The predefined attributes of namespace `USERENV` are listed [Table 4-1](#) on page 4-94, along with the lengths of their return strings.

#### See Also:

- *Oracle8i Application Developer's Guide - Fundamentals* for information on using the application context feature in your application development.
- "[CREATE CONTEXT](#)" on page 7-261 for information on creating user-defined context namespaces.
- *Oracle8i Supplied PL/SQL Packages Reference* for information on the `DBMS_SESSION.SET_CONTEXT` procedure.

### Example 1

The following statement returns the name of the user who logged onto the database:

```

SELECT SYS_CONTEXT ( 'USERENV' , 'SESSION_USER' )
       FROM DUAL;

```

```
SYS_CONTEXT ('USERENV', 'SESSION_USER')
```

```
-----  
SCOTT
```

## Example 2

The following example returns the group number that was set as the value for the attribute `GROUP_NO` in the PL/SQL package that was associated with the context `HR_APPS` when `HR_APPS` was created:

```
SELECT SYS_CONTEXT ('hr_apps', 'group_no') "User Group"  
FROM DUAL;
```

```
User Group
```

```
-----  
Sales
```

**Table 4–1** *Predefined Attributes of Namespace USERENV*

Attribute	Return Value	Return Length (bytes)
TERMINAL	The operating system identifier for the client of the current session. In distributed SQL statements, this option returns the identifier for your local session. In a distributed environment, this is supported only for remote <code>SELECTs</code> , not for remote <code>INSERTs</code> , <code>UPDATEs</code> , or <code>DELETEs</code> . (The return length of this parameter may vary by operating system.)	10
LANGUAGE	The language and territory currently used by your session, along with the database character set, in this form:  <code>language_territory.characterset</code>	52
LANG	The ISO abbreviation for the language name, a shorter form than the existing <code>'LANGUAGE'</code> parameter.	62
SESSIONID	The auditing session identifier. You cannot use this option in distributed SQL statements.	30
INSTANCE	The instance identification number of the current instance.	30



**Table 4–1 Predefined Attributes of Namespace USERENV**

<b>Attribute</b>	<b>Return Value</b>	<b>Return Length (bytes)</b>
ENTRYID	The available auditing entry identifier. You cannot use this option in distributed SQL statements. To use this keyword in USERENV, the initialization parameter AUDIT_TRAIL must be set to TRUE.	30
ISDBA	TRUE if you currently have the DBA role enabled and FALSE if you do not.	30
CLIENT_INFO	Returns up to 64 bytes of user session information that can be stored by an application using the DBMS_APPLICATION_INFO package.	64
NLS_TERRITORY	The territory of the current session.	62
NLS_CURRENCY	The currency of the current session.	62
NLS_CALENDAR	The current calendar of the current session.	62
NLS_DATE_FORMAT	The date format for the session.	62
NLS_DATE_LANGUAGE	The language used for expressing dates.	62
NLS_SORT	BINARY or the linguistic sort basis.	62
CURRENT_USER	The name of the user whose privilege the current session is under.	30
CURRENT_USERID	User ID of the user whose privilege the current session is under	30
SESSION_USER	Database user name by which the current user is authenticated. This value remains the same throughout the duration of the session.	30
SESSION_USERID	Identifier of the database user name by which the current user is authenticated.	30
CURRENT_SCHEMA	Name of the default schema being used in the current schema. This value can be changed during the session with an ALTER SESSION SET CURRENT_SCHEMA statement.	30
CURRENT_SCHEMAID	Identifier of the default schema being used in the current session.	30
PROXY_USER	Name of the database user who opened the current session on behalf of SESSION_USER.	30

**Table 4–1 Predefined Attributes of Namespace USERENV**

<b>Attribute</b>	<b>Return Value</b>	<b>Return Length (bytes)</b>
PROXY_USERID	Identifier of the database user who opened the current session on behalf of SESSION_USER.	30
DB_DOMAIN	Domain of the database as specified in the DB_DOMAIN initialization parameter.	256
DB_NAME	Name of the database as specified in the DB_NAME initialization parameter	30
HOST	Name of the host machine from which the client has connected.	54
OS_USER	Operating system username of the client process that initiated the database session	30
EXTERNAL_NAME	External name of the database user. For SSL authenticated sessions using v.503 certificates, this field returns the distinguished name (DN) stored in the user certificate.	256
IP_ADDRESS	IP address of the machine from which the client is connected.	30
NETWORK_PROTOCOL	Network protocol being used for communication, as specified in the 'PROTOCOL= <i>protocol</i> ' portion of the connect string.	256
BG_JOB_ID	Job ID of the current session if it was established by an Oracle background process. Null if the session was not established by a background process.	30
FG_JOB_ID	Job ID of the current session if it was established by a client foreground process. Null if the session was not established by a foreground process.	30
AUTHENTICATION_TYPE	How the user was authenticated: <ul style="list-style-type: none"> <li>■ DATABASE: username/password authentication</li> <li>■ OS: operating system external user authentication</li> <li>■ NETWORK: network protocol or ANO authentication</li> <li>■ PROXY: OCI proxy connection authentication</li> </ul>	30
AUTHENTICATION_DATA	Data being used to authenticate the login user. For X.503 certificate authenticated sessions, this field returns the context of the certificate in HEX2 format.	256

**Table 4–1** Predefined Attributes of Namespace USERENV

Attribute	Return Value	Return Length (bytes)
	<p><b>Note:</b> You can change the return value of the AUTHENTICATION_DATA attribute using the <i>length</i> parameter of the syntax. Values of up to 4000 are accepted. This is the only attribute of USERENV for which Oracle implements such a change.</p>	

## SYS\_GUID

### Syntax

```
→ SYS_GUID → ( ) →
```

### Purpose

`SYS_GUID` generates and returns a globally unique identifier (RAW value) made up of 16 bytes. On most platforms, the generated identifier consists of a host identifier and a process or thread identifier of the process or thread invoking the function, and a nonrepeating value (sequence of bytes) for that process or thread.

### Example

The second line of this example returns the 32-character hexadecimal representation of the 16-byte raw value of the global unique identifier.

```
INSERT INTO my_table VALUES ('BOB', SYS_GUID());
```

```
SELECT SYS_GUID() FROM DUAL;
```

```
SYS_GUID()
```

```
-----  
54FA6C5E19733A05E03400400B40DCB1
```

## SYSDATE

### Syntax

→ SYSDATE →

### Purpose

SYSDATE returns the current date and time. Requires no arguments. In distributed SQL statements, this function returns the date and time on your local database. You cannot use this function in the condition of a CHECK constraint.

### Example

```
SELECT TO_CHAR
      (SYSDATE, 'MM-DD-YYYY HH24:MI:SS') "NOW"
FROM DUAL;
```

```
NOW
-----
10-29-1999 20:27:11
```

## TAN

### Syntax

→ TAN ( *n* ) →

### Purpose

TAN returns the tangent of *n* (an angle expressed in radians).

### Example

```
SELECT TAN(135 * 3.14159265359/180)
      "Tangent of 135 degrees" FROM DUAL;
```

```
Tangent of 135 degrees
-----
- 1
```

## TANH

### Syntax



### Purpose

TANH returns the hyperbolic tangent of *n*.


### Example

```
SELECT TANH(.5) "Hyperbolic tangent of .5"
       FROM DUAL;
```

```
Hyperbolic tangent of .5
-----
                .462117157
```

## TO\_CHAR (date conversion)

### Syntax



### Purpose

TO\_CHAR converts *d* of DATE datatype to a value of VARCHAR2 datatype in the format specified by the date format *fmt*. If you omit *fmt*, *d* is converted to a VARCHAR2 value in the default date format. For information on date formats, see ["Format Models"](#) on page 2-38.

The '*nlsparams*' specifies the language in which month and day names and abbreviations are returned. This argument can have this form:

```
'NLS_DATE_LANGUAGE = language'
```


If you omit *nlsparams*, this function uses the default date language for your session.

**Example**

```
SELECT TO_CHAR(HIREDATE, 'Month DD, YYYY')
       "New date format" FROM emp
       WHERE ename = 'BLAKE';
```

New date format

```
-----
May          01, 1981
```

**TO\_CHAR (number conversion)****Syntax****Purpose**

TO\_CHAR converts *n* of NUMBER datatype to a value of VARCHAR2 datatype, using the optional number format *fmt*. If you omit *fmt*, *n* is converted to a VARCHAR2 value exactly long enough to hold its significant digits. For information on number formats, see ["Format Models"](#) on page 2-38.

The '*nlsparams*' specifies these characters that are returned by number format elements:

- decimal character
- group separator
- local currency symbol
- international currency symbol

This argument can have this form:

```
'NLS_NUMERIC_CHARACTERS = 'dg'
  NLS_CURRENCY = 'text'
  NLS_ISO_CURRENCY = territory '
```

The characters *d* and *g* represent the decimal character and group separator, respectively. They must be different single-byte characters. Note that within the

quoted string, you must use two single quotation marks around the parameter values. Ten characters are available for the currency symbol.

If you omit *nlsparms* or any one of the parameters, this function uses the default parameter values for your session.

### Example 1

In this example, the output is blank padded to the left of the currency symbol.

```
SELECT TO_CHAR(-10000, 'L99G999D99MI') "Amount"
       FROM DUAL;
```

```
Amount
-----
  $10,000.00-
```

### Example 2

```
SELECT TO_CHAR(-10000, 'L99G999D99MI',
       'NLS_NUMERIC_CHARACTERS = ','.'
       NLS_CURRENCY = 'AusDollars') "Amount"
       FROM DUAL;
```

```
Amount
-----
AusDollars10.000,00-
```

---



---

**Note:** In the optional number format *fmt*, L designates local currency symbol and MI designates a trailing minus sign. See [Table 2-7](#) on page 2-41 for a complete listing of number format elements.


---



---

## TO\_DATE

### Syntax



## Purpose

`TO_DATE` converts *char* of CHAR or VARCHAR2 datatype to a value of DATE datatype. The *fmt* is a date format specifying the format of *char*. If you omit *fmt*, *char* must be in the default date format. If *fmt* is 'J', for Julian, then *char* must be an integer.

The '*nlsparams*' has the same purpose in this function as in the `TO_CHAR` function for date conversion.

Do not use the `TO_DATE` function with a DATE value for the *char* argument. The returned DATE value can have a different century value than the original *char*, depending on *fmt* or the default date format.

**See Also:** ["Date Format Models"](#) on page 2-45.

## Example

```
INSERT INTO bonus (bonus_date)
  SELECT TO_DATE(
    'January 15, 1989, 11:00 A.M.',
    'Month dd, YYYY, HH:MI A.M.',
    'NLS_DATE_LANGUAGE = American')
  FROM DUAL;
```

# TO\_LOB

## Syntax

→ TO\_LOB → ( → long\_column → ) →

## Purpose

`TO_LOB` converts LONG or LONG RAW values in the column *long\_column* to LOB values. You can apply this function only to a LONG or LONG RAW column, and only in the SELECT list of a subquery in an INSERT statement (see ["INSERT"](#) on page 7-539).

Before using this function, you must create a LOB column to receive the converted LONG values. To convert LONGs, create a CLOB column. To convert LONG RAWs, create a BLOB column.

## Example

Given the following tables:



```
CREATE TABLE long_table (n NUMBER, long_col LONG);
CREATE TABLE lob_table (n NUMBER, lob_col CLOB);
```

use this function to convert LONG to LOB values as follows:

```
INSERT INTO lob_table
  SELECT n, TO_LOB(long_col) FROM long_table;
```

## TO\_MULTI\_BYTE

### Syntax

```
→ [TO_MULTI_BYTE] → ( ( char ) ) →
```

### Purpose

TO\_MULTI\_BYTE returns *char* with all of its single-byte characters converted to their corresponding multibyte characters. Any single-byte characters in *char* that have no multibyte equivalents appear in the output string as single-byte characters. This function is useful only if your database character set contains both single-byte and multibyte characters.

## TO\_NUMBER

### Syntax

```
→ [TO_NUMBER] → ( ( char [, fmt [, nlsparam]] ) ) →
```

### Purpose

TO\_NUMBER converts *char*, a value of CHAR or VARCHAR2 datatype containing a number in the format specified by the optional format model *fmt*, to a value of NUMBER datatype.

**Example 1**

```
UPDATE emp SET sal = sal +
    TO_NUMBER('100.00', '9G999D99')
WHERE ename = 'BLAKE';
```

The *'nlsparams'* string in this function has the same purpose as it does in the TO\_CHAR function for number conversions.

**See Also:** ["TO\\_CHAR \(number conversion\)"](#) on page 4-100.

**Example 2**

```
SELECT TO_NUMBER('-AusDollars100','L9G999D99',
    ' NLS_NUMERIC_CHARACTERS = ','.'
    NLS_CURRENCY           = 'AusDollars'
) "Amount"
FROM DUAL;
```

```
Amount
-----
-100
```

## TO\_SINGLE\_BYTE

**Syntax**

→ TO\_SINGLE\_BYTE ( ( char ) ) →

**Purpose**

TO\_SINGLE\_BYTE returns *char* with all of its multibyte characters converted to their corresponding single-byte characters. Any multibyte characters in *char* that have no single-byte equivalents appear in the output as multibyte characters. This function is useful only if your database character set contains both single-byte and multibyte characters.

# TRANSLATE

## Syntax

```
→ [TRANSLATE] → ( → ' → char → ' → , → ' → from → ' → , → ' → to → ' → ) →
```

## Purpose

TRANSLATE returns *char* with all occurrences of each character in *from* replaced by its corresponding character in *to*. Characters in *char* that are not in *from* are not replaced. The argument *from* can contain more characters than *to*. In this case, the extra characters at the end of *from* have no corresponding characters in *to*. If these extra characters appear in *char*, they are removed from the return value. You cannot use an empty string for *to* to remove all characters in *from* from the return value. Oracle interprets the empty string as null, and if this function has a null argument, it returns null.

## Example 1

The following statement translates a license number. All letters 'ABC...Z' are translated to 'X' and all digits '012 . . . 9' are translated to '9':

```
SELECT TRANSLATE('2KRW229',
'0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ',
'9999999999XXXXXXXXXXXXXXXXXXXXXXXXXXXX') "License"
      FROM DUAL;
```

```
License
-----
9XXX999
```

## Example 2


The following statement returns a license number with the characters removed and the digits remaining:

```
SELECT TRANSLATE('2KRW229',
'0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZ', '0123456789')
"Translate example"
      FROM DUAL;
```

```
Translate example
-----
2229
```

## TRANSLATE ... USING

### Syntax



### Purpose

`TRANSLATE ... USING` converts *text* into the character set specified for conversions between the database character set and the national character set.

The *text* argument is the expression to be converted.

Specifying the `USING CHAR_CS` argument converts *text* into the database character set. The output datatype is `VARCHAR2`.

Specifying the `USING NCHAR_CS` argument converts *text* into the national character set. The output datatype is `NVARCHAR2`.

This function is similar to the Oracle `CONVERT` function, but must be used instead of `CONVERT` if either the input or the output datatype is being used as `NCHAR` or `NVARCHAR2`.

### Example 1

```

CREATE TABLE t1 (char_col CHAR(20),
                 nchar_col nchar(20));

INSERT INTO t1
VALUES ('Hi', N'Bye');

SELECT * FROM t1;

```

CHAR_COL	NCHAR_COL
-----	-----
Hi	Bye

### Example 2

```

UPDATE t1 SET
  nchar_col = TRANSLATE(char_col USING NCHAR_CS);
UPDATE t1 SET
  char_col = TRANSLATE(nchar_col USING CHAR_CS);
SELECT * FROM t1;

```

CHAR_COL	NCHAR_COL
-----	-----
Hi	Hi

### Example 3

```


UPDATE t1 SET
  nchar_col = TRANSLATE('deo' USING NCHAR_CS);
UPDATE t1 SET
  char_col = TRANSLATE(N'deo' USING CHAR_CS);
SELECT * FROM t1;

```

CHAR_COL	NCHAR_COL
-----	-----
deo	deo

## TRIM

### Syntax



### Purpose

TRIM enables you to trim leading or trailing characters (or both) from a character string. If *trim\_character* or *trim\_source* is a character literal, you must enclose it in single quotes.

- If you specify **LEADING**, Oracle removes any leading characters equal to *trim\_character*.
- If you specify **TRAILING**, Oracle removes any trailing characters equal to *trim\_character*.
- If you specify **BOTH** or none of the three, Oracle removes leading and trailing characters equal to *trim\_character*.

- If you do not specify *trim\_character*, the default value is a blank space.
- If you specify only *trim\_source*, Oracle removes leading and trailing blank spaces.
- The function returns a value with datatype VARCHAR2. The maximum length of the value is the length of *trim\_source*.
- If either *trim\_source* or *trim\_character* is a null value, then the TRIM function returns a null value.

This example trims leading and trailing zeroes from a number:


### Example

```
SELECT TRIM (0 FROM 0009872348900) "TRIM Example"
FROM DUAL;
```

```
TRIM example
-----
      98723489
```

## TRUNC (Number Function)

### Syntax



### Purpose

TRUNC returns *n* truncated to *m* decimal places. If *m* is omitted, *n* is truncated to 0 places. *m* can be negative to truncate (make zero) *m* digits left of the decimal point.

### Example

```
SELECT TRUNC(15.79,1) "Truncate" FROM DUAL;
```

```
Truncate
-----
      15.7
```


```
SELECT TRUNC(15.79,-1) "Truncate" FROM DUAL;
```

```

Truncate
-----
          10
    
```

## TRUNC (Date Function)

### Syntax



### Purpose

TRUNC returns *d* with the time portion of the day truncated to the unit specified by the format model *fmt*. If you omit *fmt*, *d* is truncated to the nearest day. See ["ROUND and TRUNC Date Functions"](#) on page 4-117 for the permitted format models to use in *fmt*.

### Example

```

SELECT TRUNC(TO_DATE('27-OCT-92','DD-MON-YY'),'YEAR')
       "New Year" FROM DUAL;
    
```

```

New Year
-----
01-JAN-92
    
```

## UID

### Syntax



### Purpose

UID returns an integer that uniquely identifies the session user (the user who logged on).

**Example**

```
SELECT UID FROM DUAL;
```

```
      UID  
-----  
      19
```

## UPPER

**Syntax**

→ UPPER → ( → char → ) →

**Purpose**

UPPER returns *char*, with all letters uppercase. The return value has the same datatype as the argument *char*.

**Example**

```
SELECT UPPER('Large') "Uppercase"  
      FROM DUAL;
```

```
Uppercase  
-----  
LARGE
```

## USER

**Syntax**

→ USER →

**Purpose**

USER returns the name of the session user (the user who logged on) with the datatype VARCHAR2. Oracle compares values of this function with blank-padded comparison semantics.



In a distributed SQL statement, the `UID` and `USER` functions identify the user on your local database. You cannot use these functions in the condition of a `CHECK` constraint.


### Example

```
SELECT USER, UID FROM DUAL;
```

USER	UID
-----	-----
SCOTT	19

## USERENV

### Syntax



```
→ [USERENV] ( ( ) option ( ) ) →
```

### Purpose

`USERENV` returns information of `VARCHAR2` datatype about the current session. This information can be useful for writing an application-specific audit trail table or for determining the language-specific characters currently used by your session. You cannot use `USERENV` in the condition of a `CHECK` constraint. The argument *option* can have any of these values:

- |             |  |
|-------------|--|
| 'ISDBA'     | returns 'TRUE' if you currently have the ISDBA role enabled and 'FALSE' if you do not.   |
| 'LANGUAGE'  | returns the language and territory currently used by your session along with the database character set in this form:<br><br>language_territory.characterset   |
| 'TERMINAL'  | returns the operating system identifier for your current session's terminal. In distributed SQL statements, this option returns the identifier for your local session. In a distributed environment, this is supported only for remote <code>SELECTs</code> , not for remote <code>INSERTs</code> , <code>UPDATEs</code> , or <code>DELETEs</code> . |
| 'SESSIONID' | returns your auditing session identifier. You cannot use this option in distributed SQL statements.  |

'ENTRYID'	returns available auditing entry identifier. You cannot use this option in distributed SQL statements. To use this keyword in USERENV, the initialization parameter AUDIT_TRAIL must be set to TRUE.
'LANG'	returns the ISO abbreviation for the language name, a shorter form than the existing 'LANGUAGE' parameter.
'INSTANCE'	returns the instance identification number of the current instance.
'CLIENT_INFO'	returns up to 64 bytes of user session information that can be stored by an application using the DBMS_APPLICATION_INFO package.

**CAUTION:** Some commercial applications may be using this context value. Check the applicable documentation for those applications to determine what restrictions they may impose on use of this context area.

Oracle recommends that you use the application context feature or the SYS\_CONTEXT function with the USERENV option. These alternatives are more secure and flexible.

**See Also:**

- *Oracle8i Concepts*. for information on application context.
- "[CREATE CONTEXT](#)" on page 7-261 and "[SYS\\_CONTEXT](#)" on page 4-93.


**Example**

```
SELECT USERENV('LANGUAGE') "Language" FROM DUAL;
```

```
Language
-----
AMERICAN_AMERICA.WE8DEC
```

**VALUE**

**Syntax**



## Purpose

In a SQL statement, `VALUE` takes as its argument a correlation variable (table alias) associated with a row of an object table and returns object instances stored in the object table. The type of the object instances is the same type as the object table.

## Example


```
CREATE TYPE emp_type AS OBJECT
  (eno NUMBER, ename VARCHAR2(20), salary NUMBER);
CREATE TABLE emp_table OF emp_type
  (primary key (eno, ename));
INSERT INTO emp_table VALUES (10, 'jack', 50000);
SELECT VALUE(e) FROM emp_table e;
```

```
VALUE(E)(ENO, ENAME, SALARY)
```

```
-----
EMP_TYPE(10, 'jack', 50000)
```

# VAR\_POP

## Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

## Purpose

`VAR_POP` returns the population variance of a set of numbers after discarding the nulls in this set. You can use it as both an aggregate and analytic function.

The `expr` is a number expression, and the function returns a value of type `NUMBER`. If the function is applied to an empty set, it returns null. The function makes the following calculation:

$$\frac{(\text{SUM}(\text{expr}^2) - \text{SUM}(\text{expr})^2 / \text{COUNT}(\text{expr}))}{\text{COUNT}(\text{expr})}$$

**See Also:** ["Aggregate Functions"](#) on page 4-6

### Aggregate Example

The following example returns the population variance of the salaries in the EMP table:

```
SELECT VAR_POP(sal) FROM emp;
```

```
VAR_POP(SAL)
-----
 1298434.31
```

### Analytic Example


The following example calculates the cumulative population and sample variances of the monthly sales in 1998:

```
SELECT s_month, VAR_POP(SUM(s_amount)) OVER (ORDER BY s_month),
       VAR_SAMP(SUM(s_amount)) OVER (ORDER BY s_month)
FROM sales
WHERE s_year =1998
GROUP BY s_month;
```

```
S_MONTH      VAR_POP(SU  VAR_SAMP(S
-----
          6          0
          8  440588496  881176992
         11  538819892  808229838
```

## VAR\_SAMP

### Syntax



For information on syntax and semantics, see "[Analytic Functions](#)" on page 4-7.

### Purpose

VAR\_SAMP returns the sample variance of a set of numbers after discarding the nulls in this set. You can use it as both an aggregate and analytic function.

The *expr* is a number expression, and the function returns a value of type NUMBER. If the function is applied to an empty set, it returns null. The function makes the following calculation:

$$(\text{SUM}(\text{expr}^2) - \text{SUM}(\text{expr})^2 / \text{COUNT}(\text{expr})) / (\text{COUNT}(\text{expr}) - 1)$$

This function is similar to VARIANCE, except that given an input set of one element, VARIANCE returns 0 and VAR\_SAMP returns null.

**See Also:** ["Aggregate Functions"](#) on page 4-6

### Aggregate Example

The following example returns the sample variance of the salaries in the EMP table.

```
SELECT VAR_SAMP(sal) FROM emp;
```


```
VAR_SAMP(SAL)
-----
      1398313.87
```

### Analytic Example

See the example at ["VAR\\_POP"](#) on page 4-113.

## VARIANCE

### Syntax



For information on syntax and semantics, see ["Analytic Functions"](#) on page 4-7.

### Purpose

VARIANCE returns variance of *expr*. You can use it as an aggregate or analytic function.

Oracle calculates the variance of *expr* as follows:

- 0 if the number of rows in *expr* = 1

- VAR\_SAMP if the number of rows in *expr* > 1

### Aggregate Example

The following example calculates the variance of all salaries in the EMP table:

```
SELECT VARIANCE(sal) "Variance"  
FROM emp;
```

```
Variance  
-----  
1389313.87
```

### Analytic Example

The query returns the cumulative variance of salary values in Department 30 ordered by hiredate.

```
SELECT ename, sal, VARIANCE(sal) OVER (ORDER BY hiredate)  
FROM emp  
WHERE deptno=30;
```

ENAME	SAL	VARIANCE(S
ALLEN	1600	0
WARD	1250	61250
BLAKE	2850	707500
TURNER	1500	511666.667
MARTIN	1250	444250
JAMES	950	446666.667

## VSIZE

### Syntax

→ VSIZE → ( → expr → ) →

### Purpose

VSIZE returns the number of bytes in the internal representation of *expr*. If *expr* is null, this function returns null.

**Example**

```
SELECT ename, VSIZE (ename) "BYTES"
   FROM emp
  WHERE deptno = 10;
```

```
ENAME          BYTES
-----
CLARK          5
KING           4
MILLER        6
```

**ROUND and TRUNC Date Functions**

[Table 4–2](#) lists the format models you can use with the ROUND and TRUNC date functions and the units to which they round and truncate dates. The default model, 'DD', returns the date rounded or truncated to the day with a time of midnight.

**Table 4–2** *Date Format Models for the ROUND and TRUNC Date Functions*

Format Model	Rounding or Truncating Unit
CC SCC	One greater than the first two digits of a four-digit year.
SYYYY YYYY YEAR SYEAR YYY YY Y	Year (rounds up on July 1)
IYYY IY IY I	ISO Year
Q	Quarter (rounds up on the sixteenth day of the second month of the quarter)
MONTH MON MM RM	Month (rounds up on the sixteenth day)
WW	Same day of the week as the first day of the year.

**Table 4–2 (Cont.) Date Format Models for the ROUND and TRUNC Date Functions**

<b>Format Model</b>	<b>Rounding or Truncating Unit</b>
IW	Same day of the week as the first day of the ISO year.
W	Same day of the week as the first day of the month.
DDD DD J	Day
DAY DY D	Starting day of the week
HH HH12 HH24	Hour
MI	Minute

The starting day of the week used by the format models DAY, DY, and D is specified implicitly by the initialization parameter `NLS_TERRITORY`.

**See Also:** *Oracle8i Reference* and *Oracle8i National Language Support Guide* for information on this parameter.

## User-Defined Functions

You can write user-defined functions in PL/SQL or Java to provide functionality that is not available in SQL or SQL functions. User functions can appear in a SQL statement anywhere SQL functions can appear, that is, wherever an expression can occur.

For example, user functions can be used in the following:

- The select list of a `SELECT` statement
- The condition of a `WHERE` clause
- `CONNECT BY`, `START WITH`, `ORDER BY`, and `GROUP BY` clauses
- The `VALUES` clause of an `INSERT` statement
- The `SET` clause of an `UPDATE` statement



**See Also:**

- ["CREATE FUNCTION"](#) on page 7-284 for information on creating functions, including restrictions on user-defined functions.
- *Oracle8i Application Developer's Guide - Fundamentals* for a complete description on the creation and use of user functions.

## Prerequisites

User functions must be created as top-level functions or declared with a package specification before they can be named within a SQL statement. Create user functions as top-level functions by using the `CREATE FUNCTION` statement described in ["CREATE FUNCTION"](#) on page 7-284. To specify packaged functions, see ["CREATE PACKAGE"](#) on page 7-344.

To use a user function in a SQL expression, you must own or have `EXECUTE` privilege on the user function. To query a view defined with a user function, you must have `SELECT` privileges on the view. No separate `EXECUTE` privileges are needed to select from the view.

## Name Precedence

Within a SQL statement, the names of database columns take precedence over the names of functions with no parameters. For example, if user `SCOTT` creates the following two objects in his own schema:

```
CREATE TABLE emp(new_sal NUMBER, ...);
CREATE FUNCTION new_sal RETURN NUMBER IS BEGIN ... END;
```

then in the following two statements, the reference to `NEW_SAL` refers to the column `EMP.NEW_SAL`:

```
SELECT new_sal FROM emp;
SELECT emp.new_sal FROM emp;
```

To access the function `NEW_SAL`, you would enter:

```
SELECT scott.new_sal FROM emp;
```

Here are some sample calls to user functions that are allowed in SQL expressions:

```
circle_area (radius)
payroll.tax_rate (empno)
scott.payroll.tax_rate (dependent, empno)@ny
```

**Example** To call the `TAX_RATE` user function from schema `SCOTT`, execute it against the `SS_NO` and `SAL` columns in `TAX_TABLE`, and place the results in the variable `INCOME_TAX`, specify the following:

```
SELECT scott.tax_rate (ss_no, sal)
       INTO income_tax
       FROM tax_table
       WHERE ss_no = tax_id;
```

### Naming Conventions

If only one of the optional schema or package names is given, the first identifier can be either a schema name or a package name. For example, to determine whether `PAYROLL` in the reference `PAYROLL.TAX_RATE` is a schema or package name, Oracle proceeds as follows:

1. Check for the `PAYROLL` package in the current schema.
2. If a `PAYROLL` package is not found, look for a schema name `PAYROLL` that contains a top-level `TAX_RATE` function. If no such function is found, return an error.
3. If the `PAYROLL` package is found in the current schema, look for a `TAX_RATE` function in the `PAYROLL` package. If no such function is found, return an error.

You can also refer to a stored top-level function using any synonym that you have defined for it.

---

# Expressions, Conditions, and Queries

*The ideal condition would be, I admit, that men should be right by instinct.*

Sophocles, *Oedipus Rex*

This chapter describes how to combine the values, operators, and functions described in earlier chapters evaluate to a value. Topics include:

- [Expressions](#)
- [Conditions](#)
- [Queries and Subqueries](#)

## Expressions

An *expression* is a combination of one or more values, operators, and SQL functions that evaluate to a value. An expression generally assumes the datatype of its components.

This simple expression evaluates to 4 and has datatype NUMBER (the same datatype as its components):

```
2*2
```

The following expression is an example of a more complex expression that uses both functions and operators. The expression adds seven days to the current date, removes the time component from the sum, and converts the result to CHAR datatype:

```
TO_CHAR ( TRUNC ( SYSDATE+7 ) )
```

You can use expressions in:

- The select list of the SELECT statement

- A condition of the WHERE clause and HAVING clause
- The CONNECT BY, START WITH, and ORDER BY clauses
- The VALUES clause of the INSERT statement
- The SET clause of the UPDATE statement

For example, you could use an expression in place of the quoted string 'smith' in this UPDATE statement SET clause:

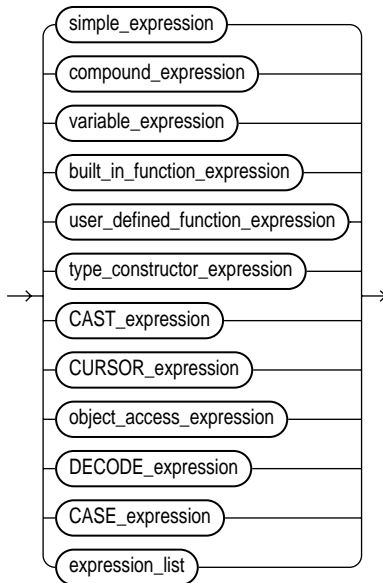
```
SET ename = 'smith';
```

This SET clause has the expression LOWER(ename) instead of the quoted string 'smith':

```
SET ename = LOWER(ename);
```

Expressions have several forms, as shown in the following syntax:

**expr::=**



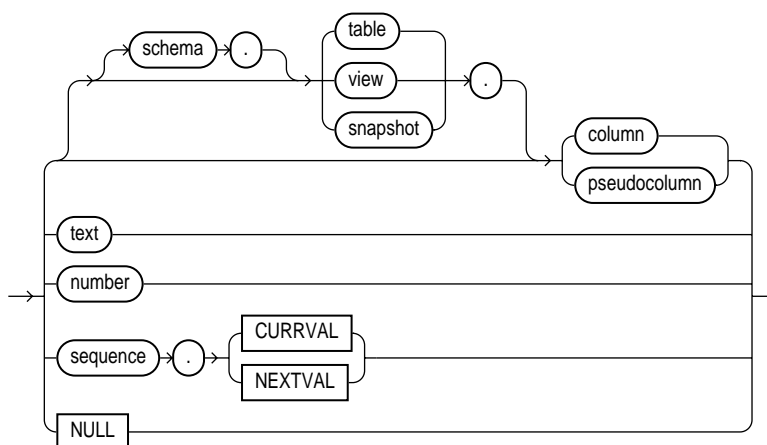
Oracle does not accept all forms of expressions in all parts of all SQL statements. You must use appropriate expression notation whenever *expr* appears in conditions, SQL functions, or SQL statements in other parts of this reference. The description of

each statement in [Chapter 7, "SQL Statements"](#), documents the restrictions on the expressions in the statement. The sections that follow describe and provide examples of the various forms of expressions.

## Simple Expressions

A simple expression specifies column, pseudocolumn, constant, sequence number, or null.

**simple\_expression::=**



In addition to the schema of a user, *schema* can also be "PUBLIC" (double quotation marks required), in which case it must qualify a public synonym for a table, view, or materialized view. Qualifying a public synonym with "PUBLIC" is supported only in data manipulation language (DML) statements, not data definition language (DDL) statements.

The *pseudocolumn* can be either LEVEL, ROWID, or ROWNUM. You can use a pseudocolumn only with a table, not with a view or materialized view. NCHAR and NVARCHAR2 are not valid pseudocolumn datatypes.

**See Also:** ["Pseudocolumns"](#) on page 2-56 for more information on pseudocolumns.

Some valid simple expressions are:

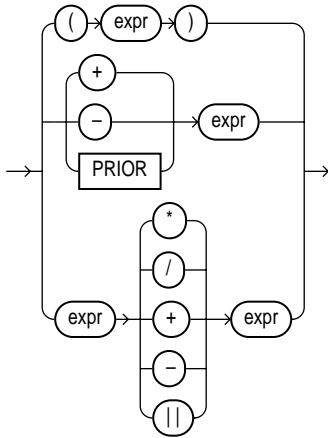
emp.ename

```
'this is a text string'
10
N'this is an NCHAR string'
```

## Compound Expressions

A compound expression specifies a combination of other expressions.

**compound\_expression::=**



Note that some combinations of functions are inappropriate and are rejected. For example, the `LENGTH` function is inappropriate within an aggregate function.


Some valid compound expressions are:

```
('CLARK' || 'SMITH')
LENGTH('MOOSE') * 57
SQRT(144) + 72
my_fun(TO_CHAR(sysdate, 'DD-MMM-YY'))
```

## Variable Expressions

A variable expression specifies a host variable with an optional indicator variable. Note that this form of expression can appear only in embedded SQL statements or SQL statements processed in an Oracle Call Interface (OCI) program.

**variable\_expression::=**




Some valid variable expressions are:

```
:employee_name INDICATOR :employee_name_indicator_var
:department_location
```

## Built-In Function Expressions

A built-in function expression specifies a call to a single-row SQL function.

**built\_in\_function\_expression::=**



Some valid built-in function expressions are:

```
LENGTH( ' BLAKE' )
ROUND( 1234.567*43 )
SYSDATE
```


For information on built-in functions, see ["SQL Functions"](#) on page 4-2. See also ["Aggregate Functions"](#) on page 4-6.

## Function Expressions

A function expression specifies a call to

- A SQL built-in function (see [Chapter 4, "Functions"](#))
- A function in an Oracle-supplied package (see *Oracle8i Supplied PL/SQL Packages Reference*).
- A function in a user-defined package or in a standalone user-defined function (see ["User-Defined Functions"](#) on page 4-118)
- A user-defined operator (see ["CREATE OPERATOR"](#) on page 7-339 and *Oracle8i Data Cartridge Developer's Guide*)

**function\_expression::=**




Some valid user-defined function expressions are:

```
circle_area(radius)
payroll.tax_rate(empno)
scott.payrol.tax_rate(dependents, empno)@ny
DBMS_LOB.getlength(column_name)
```

## Type Constructor Expressions

A type constructor expression specifies a call to a type constructor. The *argument* to the type constructor is any expression or subquery.

**type\_constructor\_expression::=**



If *type\_name* is an **object type**, then the argument list must be an ordered list, where the first argument is a value whose type matches the first attribute of the object type, the second argument is a value whose type matches the second attribute of the object type, and so on. The total number of arguments to the constructor must match the total number of attributes of the object type.

If *type\_name* is a **varray or nested table type**, then the argument list can contain zero or more arguments. Zero arguments implies construction of an empty collection. Otherwise, each argument corresponds to an element value whose type is the element type of the collection type.

If *type\_name* is an **object type, a varray, or a nested table type**, the maximum number of arguments it can contain is 1000 minus some overhead.



**Expression Example** This example shows the use of an expression in the call to a type constructor.

```
CREATE TYPE address_t AS OBJECT
  (no NUMBER, street CHAR(31), city CHAR(21), state CHAR(3), zip NUMBER);
CREATE TYPE address_book_t AS TABLE OF address_t;
DECLARE
  /* Object Type variable initialized via Object Type Constructor */
  myaddr address_t = address_t(500, 'Oracle Parkway', 'Redwood Shores', 'CA', 94065);
  /* nested table variable initialized to an empty table via a constructor*/
  alladdr address_book_t = address_book_t();
BEGIN
  /* below is an example of a nested table constructor with two elements
     specified, where each element is specified as an object type constructor. */
  insert into employee values (666999, address_book_t(address_t(500,
    'Oracle Parkway', 'Redwood Shores', 'CA', 94065), address_t(400,
    'Mission Street', 'Fremont', 'CA', 94555)));
END;
```


**Subquery Example** This example illustrates the use of a subquery in the call to the type constructor.

```
CREATE TYPE employee AS OBJECT (
  empno NUMBER,
  ename VARCHAR2(20));
CREATE TABLE empdbl OF EMPLOYEE;
INSERT INTO empdbl VALUES(7377, 'JOHN');
CREATE TYPE project AS OBJECT (
  pname VARCHAR2(25),
  empref REF employee);
CREATE TABLE depttbl (dno number, proj project);
INSERT INTO depttbl values(10, project('SQL Extensions',
  (SELECT REF(p) FROM empdbl p
   WHERE ename='JOHN')));
```

## CAST Expressions

A CAST expression converts one built-in datatype or collection-typed value into another built-in datatype or collection-typed value.

**CAST\_expression::=**



CAST allows you to convert built-in datatypes or collection-typed values of one type into another built-in datatype or collection type. You can cast an unnamed operand (such as a date or the result set of a subquery) or a named collection (such as a varray or a nested table) into a type-compatible datatype or named collection. The *type\_name* must be the name of a built-in datatype or collection type and the *operand* must be a built-in datatype or must evaluate to a collection value.

For the operand, *expr* can be either a built-in datatype or a collection type, and *subquery* must return a single value of collection type or built-in type. `MULTISET` informs Oracle to take the result set of the subquery and return a collection value. [Table 5-1](#) shows which built-in datatypes can be cast into which other built-in datatypes. (CAST does not support LONG, LONG RAW, or any of the LOB datatypes.)

**Table 5-1 Casting Built-In Datatypes**

From/ To	CHAR, VARCHAR2	NUMBER	DATE	RAW	ROWID, UROWID	NCHAR, NVARCHAR2
CHAR, VARCHAR2	X	X	X	X	X	
NUMBER	X	X				
DATE	X		X			
RAW	X			X		
ROWID, UROWID	X				X <sup>a</sup>	
NCHAR, NVARCHAR2		X	X	X	X	X

<sup>a</sup> You cannot cast a UROWID to a ROWID if the UROWID contains the value of a ROWID of an index-organized table.

To cast a named collection type into another named collection type, the elements of both collections must be of the same type.

If the result set of *subquery* can evaluate to multiple rows, you must specify the `MULTISET` keyword. The rows resulting from the subquery form the elements of the collection value into which they are cast. Without the `MULTISET` keyword, the

subquery is treated as a scalar subquery, which is not supported in the CAST expression. In other words, scalar subqueries as arguments of the CAST operator are not valid in Oracle8i.

### Built-In Datatype Examples

```
SELECT CAST ('1997-10-22' AS DATE) FROM DUAL;
SELECT * FROM t1 WHERE CAST (ROWID AS VARCHAR2) = '01234';
```

**Collection Examples** The CAST examples that follow use the following user-defined types and tables:

```
CREATE TYPE address_t AS OBJECT
    (no NUMBER, street CHAR(31), city CHAR(21), state CHAR(2));
CREATE TYPE address_book_t AS TABLE OF address_t;
CREATE TYPE address_array_t AS VARRAY(3) OF address_t;
CREATE TABLE emp_address (empno NUMBER, no NUMBER, street CHAR(31),
    city CHAR(21), state CHAR(2));
CREATE TABLE employees (empno NUMBER, name CHAR(31));
CREATE TABLE dept (dno NUMBER, addresses address_array_t);
```

This example casts a subquery:

```
SELECT e.empno, e.name, CAST(MULTISET(SELECT ea.no, ea.street,
    ea.city, ea.state
    FROM emp_address ea
    WHERE ea.empno = e.empno)
    AS address_book_t)
    FROM employees e;
```

CAST converts a varray type column into a nested table:

```
SELECT CAST(d.addresses AS address_book_t)
    FROM dept d
    WHERE d.dno = 111;
```

The following example casts a MULTISET expression with an ORDER BY clause:

```
CREATE TABLE projects (empid NUMBER, projname VARCHAR2(10));
CREATE TABLE employees (empid NUMBER, ename VARCHAR2(10));
CREATE TYPE projname_table_type AS TABLE OF VARCHAR2(10);
```

An example of a MULTISET expression with the above schema is:

```
SELECT e.ename, CAST(MULTISET(SELECT p.projname
    FROM projects p
```

```


WHERE p.empid=e.empid
ORDER BY p.projname)
AS projname_table_type)
FROM employees e;

```

## CURSOR Expressions

A **CURSOR** expression returns a nested cursor. This form of expression is similar to the PL/SQL REF cursor.

**CURSOR\_expression::=**



A nested cursor is implicitly opened when the containing row is fetched from the parent cursor. The nested cursor is closed only when:

- The nested cursor is explicitly closed by the user
- The parent cursor is reexecuted
- The parent cursor is closed
- The parent cursor is cancelled
- An error arises during fetch on one of its parent cursors (it is closed as part of the clean-up)

**Restrictions:** The following restrictions apply to the **CURSOR** expression:

- Nested cursors can appear only in a **SELECT** statement that is not nested in any other query expression, except when it is a subquery of the **CURSOR** expression itself.
- Nested cursors can appear only in the outermost **SELECT** list of the query specification.
- Nested cursors cannot appear in views.
- You cannot perform **BIND** and **EXECUTE** operations on nested cursors.

### Example

```

SELECT d.deptno, CURSOR(SELECT e.empno, CURSOR(SELECT p.projnum,
                                             p.projname
FROM    projects p
WHERE   p.empno = e.empno)

```

```

FROM dept d
WHERE d.dno = 605;

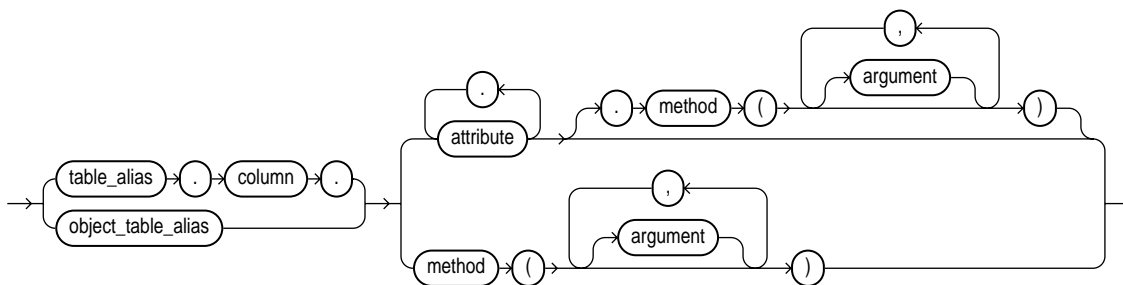
FROM TABLE(d.employees) e)

```

## Object Access Expressions

An object access expression specifies attribute reference and method invocation.

**object\_access\_expression::=**



The column parameter can be an object or REF column.

When a type's member function is invoked in the context of a SQL statement, if the SELF argument is null, Oracle returns null and the function is not invoked.

Examples in this section use the following user-defined types and tables:

```

CREATE OR REPLACE TYPE employee_t AS OBJECT
  (empid NUMBER,
   name VARCHAR2(31),
   birthdate DATE,
   MEMBER FUNCTION age RETURN NUMBER,
   PRAGMA RESTRICT_REFERENCES (age, RNPS, WNPS, WNDS)
  );

CREATE OR REPLACE TYPE BODY employee_t AS
  MEMBER FUNCTION age RETURN NUMBER IS
    var NUMBER;
  BEGIN
    var := TRUNC(MONTHS_BETWEEN(SYSDATE, birthdate) /12);
    RETURN(var);
  END;
END;

CREATE TABLE department (dno NUMBER, manager EMPLOYEE_T);

```


**Examples** The following examples update and select from the object columns and method defined above.

```
UPDATE department d
  SET d.manager.empid = 100;
SELECT d.manager.name, d.manager.age()
  FROM department d;
```

## DECODE Expressions

A `DECODE` expression uses the special `DECODE` syntax:

**DECODE\_expression::=**



To evaluate this expression, Oracle compares *expr* to each *search* value one by one. If *expr* is equal to a *search*, Oracle returns the corresponding *result*. If no match is found, Oracle returns *default*, or, if *default* is omitted, returns null. If *expr* and *search* contain character data, Oracle compares them using nonpadded comparison semantics. For information on these semantics, see the section "[Datatype Comparison Rules](#)" on page 2-32.

The *search*, *result*, and *default* values can be derived from expressions. Oracle evaluates each *search* value only before comparing it to *expr*, rather than evaluating all *search* values before comparing any of them with *expr*. Consequently, Oracle never evaluates a *search* if a previous *search* is equal to *expr*.

Oracle automatically converts *expr* and each *search* value to the datatype of the first *search* value before comparing. Oracle automatically converts the return value to the same datatype as the first *result*. If the first *result* has the datatype `CHAR` or if the first *result* is null, then Oracle converts the return value to the datatype `VARCHAR2`. For information on datatype conversion, see "[Data Conversion](#)" on page 2-36.

In a `DECODE` expression, Oracle considers two nulls to be equivalent. If *expr* is null, Oracle returns the *result* of the first *search* that is also null.

The maximum number of components in the `DECODE` expression, including *expr*, *searches*, *results*, and *default* is 255.


**Example** This expression decodes the value DEPTNO. If DEPTNO is 10, the expression evaluates to 'ACCOUNTING'; if DEPTNO is 20, it evaluates to 'RESEARCH'; etc. If DEPTNO is not 10, 20, 30, or 40, the expression returns 'NONE'.

```
DECODE (deptno,10, 'ACCOUNTING',
        20, 'RESEARCH',
        30, 'SALES',
        40, 'OPERATION',
        'NONE')
```

## CASE Expressions

CASE expressions let you use IF ... THEN ... ELSE logic in SQL statements without having to invoke procedures. The syntax is:

**CASE\_expression::=**



Oracle searches for the first WHEN ... THEN pair for which *condition* is true.

- If Oracle finds such a pair, then the result of the CASE expression is *expr1*.
- If Oracle does not find such a pair,
  - If an ELSE clause is specified, the result of the CASE expression is *expr2*.
  - If no ELSE clause is specified, the result of the CASE expression is null.

At least one occurrence of *expr1* or *expr2* must be non-null.

---

**Note:** The maximum number of arguments in a CASE expression is 255, and each WHEN ... THEN pair counts as two arguments. To avoid exceeding the limit of 128 choices, you can nest CASE expressions. That is *expr1* can itself be a CASE expression.

---


**Example** The following statement finds the average salary of all employees in the EMP table. If an employee's salary is less than \$2000, the CASE expression uses \$2000 instead.

```
SELECT AVG(CASE WHEN e.sal > 2000 THEN e.sal ELSE 2000 END) from emp e;
```

## Expression List

An expression list is a series of expressions separated by a comma. The entire series is enclosed in parentheses.

**expression\_list::=**



An expression list can contain up to 1000 expressions. Some valid expression lists are:

```
(10, 20, 40)
('SCOTT', 'BLAKE', 'TAYLOR')
(LENGTH('MOOSE') * 57, -SQRT(144) + 72, 69)
```

## Conditions

A condition specifies a combination of one or more expressions and logical operators that evaluates to either `TRUE`, `FALSE`, or unknown. You must use this syntax whenever *condition* appears in SQL statements in [Chapter 7, "SQL Statements"](#).

You can use a condition in the `WHERE` clause of these statements:

- `DELETE`
- `SELECT`
- `UPDATE`

You can use a condition in any of these clauses of the `SELECT` statement:

- `WHERE`
- `START WITH`
- `CONNECT BY`
- `HAVING`

A condition could be said to be of the "logical" datatype, although Oracle does not formally support such a datatype.

The following simple condition always evaluates to `TRUE`:



```
1 = 1
```

The following more complex condition adds the SAL value to the COMM value (substituting the value 0 for null) and determines whether the sum is greater than the number constant 2500:

```
NVL(sal, 0) + NVL(comm, 0) > 2500
```

Logical operators can combine multiple conditions into a single condition. For example, you can use the AND operator to combine two conditions:

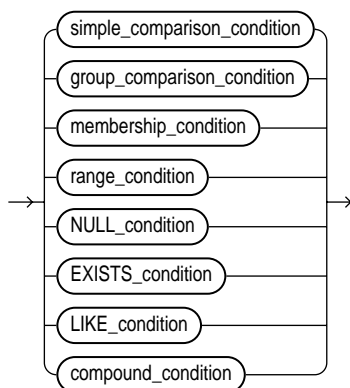
```
(1 = 1) AND (5 < 7)
```

Here are some valid conditions:

```
name = 'SMITH'
emp.deptno = dept.deptno
hiredate > '01-JAN-88'
job IN ('PRESIDENT', 'CLERK', 'ANALYST')
sal BETWEEN 500 AND 1000
comm IS NULL AND sal = 2000
```

Conditions can have several forms, as shown in the following syntax. The description of each statement in [Chapter 7, "SQL Statements"](#), documents the restrictions on the conditions in the statement. The sections that follow describe the various forms of conditions.

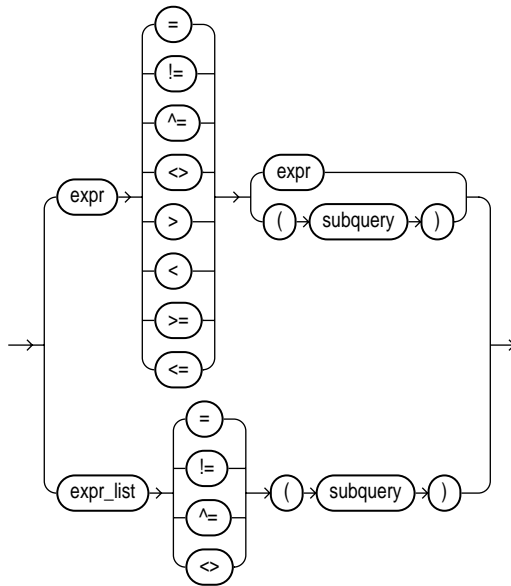
**condition::=**



## Simple Comparison Conditions

A simple comparison condition specifies a comparison with expressions or subquery results.

**simple\_comparison\_condition::=**

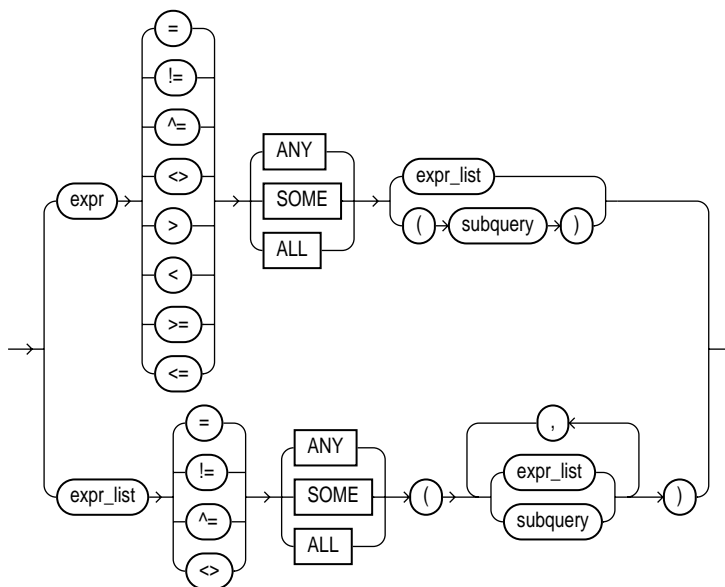


For information on comparison operators, see "[Comparison Operators](#)" on page 3-5.

## Group Comparison Conditions

A group comparison condition specifies a comparison with any or all members in a list or subquery.

**group\_comparison\_condition::=**




See "[SELECT and Subqueries](#)" on page 7-569.

## Membership Conditions

A membership condition tests for membership in a list or subquery.


**membership\_condition::=**



## Range Conditions

A range condition tests for inclusion in a range.


**range\_condition::=**



## NULL Conditions

A NULL condition tests for nulls.


**NULL\_condition::=**



## EXISTS Conditions

An EXISTS condition tests for existence of rows in a subquery.


**EXISTS\_condition::=**



## LIKE Conditions

A LIKE condition specifies a test involving pattern matching.


**LIKE\_condition::=**



## Compound Conditions

A compound condition specifies a combination of other conditions.

**compound\_condition::=**



## Queries and Subqueries

A **query** is an operation that retrieves data from one or more tables or views. In this reference, a top-level `SELECT` statement is called a **query**, and a query nested within another SQL statement is called a **subquery**.

This section describes some types of queries and subqueries and how to use them. The full syntax of all the clauses, and the semantics of the keywords and parameters, appear in "[SELECT and Subqueries](#)" on page 7-569.

### Creating Simple Queries

The list of expressions that appears after the `SELECT` keyword and before the `FROM` clause is called the *select list*. Each expression *expr* becomes the name of one column in the set of returned rows, and each *table.\** becomes a set of columns, one for each column in the table in the order they were defined when the table was created. The datatype and length of each expression is determined by the elements of the expression.

If two or more tables have some column names in common, you must qualify column names with names of tables. Otherwise, fully qualified column names are optional. However, it is always a good idea to qualify table and column references explicitly. Oracle often does less work with fully qualified table and column names.


You can use a column alias, *c\_alias*, to label the preceding expression in the select list so that the column is displayed with a new heading. The alias effectively renames the select list item for the duration of the query. The alias can be used in the `ORDER BY` clause, but not other clauses in the query.

You can use comments in a `SELECT` statement to pass instructions, or *hints*, to the Oracle optimizer. The optimizer uses hints to choose an execution plan for the statement.

**See Also:** ["Hints"](#) on page 2-63 and *Oracle8i Designing and Tuning for Performance* for more information on hints.

## Hierarchical Queries

If a table contains hierarchical data, you can select rows in a hierarchical order using the hierarchical query clause:




**START WITH** specifies the root row(s) of the hierarchy.

**CONNECT BY** specifies the relationship between parent rows and child rows of the hierarchy. Some part of *condition* must use the `PRIOR` operator to refer to the parent row. See the [PRIOR operator](#) on page 3-16.

**WHERE** restricts the rows returned by the query without affecting other rows of the hierarchy.

Oracle uses the information from the hierarchical query clause to form the hierarchy using the following steps:

1. Oracle selects the root row(s) of the hierarchy—those rows that satisfy the `START WITH condition`.
2. Oracle selects the child rows of each root row. Each child row must satisfy the condition of the `CONNECT BY condition` with respect to one of the root rows.
3. Oracle selects successive generations of child rows. Oracle first selects the children of the rows returned in step 2, and then the children of those children, and so on. Oracle always selects children by evaluating the `CONNECT BY condition` with respect to a current parent row.
4. If the query contains a `WHERE` clause, Oracle eliminates all rows from the hierarchy that do not satisfy the condition of the `WHERE` clause. Oracle evaluates this condition for each row individually, rather than removing all the children of a row that does not satisfy the condition.
5. Oracle returns the rows in the order shown in [Figure 5-1](#). In the diagram children appear below their parents.

**Figure 5–1 Hierarchical Queries**

To find the children of a parent row, Oracle evaluates the `PRIOR` expression of the `CONNECT BY` condition for the parent row and the other expression for each row in the table. Rows for which the condition is true are the children of the parent. The `CONNECT BY` condition can contain other conditions to further filter the rows selected by the query. The `CONNECT BY` condition cannot contain a subquery.

If the `CONNECT BY` condition results in a loop in the hierarchy, Oracle returns an error. A loop occurs if one row is both the parent (or grandparent or direct ancestor) and a child (or a grandchild or a direct descendent) of another row.

## Sorting Query Results

You can use the `ORDER BY` clause to order the rows selected by a query. Sorting by position is useful in the following cases:

- To order by a lengthy select list expression, you can specify its position, rather than duplicate the entire expression, in the `ORDER BY` clause.
- For compound queries (containing set operators `UNION`, `INTERSECT`, `MINUS`, or `UNION ALL`), the `ORDER BY` clause must use positions, rather than explicit expressions. Also, the `ORDER BY` clause can appear only in the last component query. The `ORDER BY` clause orders all rows returned by the entire compound query.

The mechanism by which Oracle sorts values for the `ORDER BY` clause is specified either explicitly by the `NLS_SORT` initialization parameter or implicitly by the `NLS_LANGUAGE` initialization parameter. For information on these parameters, see *Oracle8i National Language Support Guide*. You can change the sort mechanism

dynamically from one linguistic sort sequence to another using the `ALTER SESSION` statement. You can also specify a specific sort sequence for a single query by using the `NLSSORT` function with the `NLS_SORT` parameter in the `ORDER BY` clause.

## Joins

A **join** is a query that combines rows from two or more tables, views, or materialized views ("snapshots"). Oracle performs a join whenever multiple tables appear in the query's `FROM` clause. The query's select list can select any columns from any of these tables. If any two of these tables have a column name in common, you must qualify all references to these columns throughout the query with table names to avoid ambiguity.

### Join Conditions

Most join queries contain `WHERE` clause conditions that compare two columns, each from a different table. Such a condition is called a **join condition**. To execute a join, Oracle combines pairs of rows, each containing one row from each table, for which the join condition evaluates to `TRUE`. The columns in the join conditions need not also appear in the select list.

To execute a join of three or more tables, Oracle first joins two of the tables based on the join conditions comparing their columns and then joins the result to another table based on join conditions containing columns of the joined tables and the new table. Oracle continues this process until all tables are joined into the result. The optimizer determines the order in which Oracle joins tables based on the join conditions, indexes on the tables, and, in the case of the cost-based optimization approach, statistics for the tables.

In addition to join conditions, the `WHERE` clause of a join query can also contain other conditions that refer to columns of only one table. These conditions can further restrict the rows returned by the join query.

### Equijoins

An **equijoin** is a join with a join condition containing an equality operator. An equijoin combines rows that have equivalent values for the specified columns. Depending on the internal algorithm the optimizer chooses to execute the join, the total size of the columns in the equijoin condition in a single table may be limited to the size of a data block minus some overhead. The size of a data block is specified by the initialization parameter `DB_BLOCK_SIZE`. See the "[Equijoin Examples](#)" on page 7-587.



## Self Joins

A **self join** is a join of a table to itself. This table appears twice in the `FROM` clause and is followed by table aliases that qualify column names in the join condition. To perform a self join, Oracle combines and returns rows of the table that satisfy the join condition. See the ["Self Join Example"](#) on page 7-588.

## Cartesian Products

If two tables in a join query have no join condition, Oracle returns their **Cartesian product**. Oracle combines each row of one table with each row of the other. A Cartesian product always generates many rows and is rarely useful. For example, the Cartesian product of two tables, each with 100 rows, has 10,000 rows. Always include a join condition unless you specifically need a Cartesian product. If a query joins three or more tables and you do not specify a join condition for a specific pair, the optimizer may choose a join order that avoids producing an intermediate Cartesian product.

## Outer Joins

An outer join extends the result of a simple join. An **outer join** returns all rows that satisfy the join condition and those rows from one table for which no rows from the other satisfy the join condition. Such rows are not returned by a simple join. To write a query that performs an outer join of tables A and B and returns all rows from A, apply the outer join operator (+) to all columns of B in the join condition. For all rows in A that have no matching rows in B, Oracle returns `NULL` for any select list expressions containing columns of B. See the syntax for an outer join in ["SELECT and Subqueries"](#) on page 7-569.

Outer join queries are subject to the following rules and restrictions:

- The (+) operator can appear only in the `WHERE` clause or, in the context of left-correlation (that is, when specifying the `TABLE` clause) in the `FROM` clause, and can be applied only to a column of a table or view.
- If A and B are joined by multiple join conditions, you must use the (+) operator in all of these conditions. If you do not, Oracle will return only the rows resulting from a simple join, but without a warning or error to advise you that you do not have the results of an outer join.
- The (+) operator can be applied only to a column, not to an arbitrary expression. However, an arbitrary expression can contain a column marked with the (+) operator.

- A condition containing the (+) operator cannot be combined with another condition using the OR logical operator.
- A condition cannot use the IN comparison operator to compare a column marked with the (+) operator with an expression.
- A condition cannot compare any column marked with the (+) operator with a subquery.

If the WHERE clause contains a condition that compares a column from table B with a constant, the (+) operator must be applied to the column so that Oracle returns the rows from table A for which it has generated NULLs for this column. Otherwise Oracle will return only the results of a simple join.

In a query that performs outer joins of more than two pairs of tables, a single table can be the NULL-generated table for only one other table. For this reason, you cannot apply the (+) operator to columns of B in the join condition for A and B and the join condition for B and C.

## Using Subqueries

Use subqueries for the following purposes:

- To define the set of rows to be inserted into the target table of an INSERT or CREATE TABLE statement
- To define the set of rows to be included in a view or materialized view ("snapshot) in a CREATE VIEW or CREATE MATERIALIZED VIEW statement
- To define one or more values to be assigned to existing rows in an UPDATE statement
- To provide values for conditions in a WHERE clause, HAVING clause, or START WITH clause of SELECT, UPDATE, and DELETE statements
- To provide values for a specified column in an INSERT ... VALUES list
- To provide values for arguments of a type constructor or a user-defined function
- To define a table to be operated on by a containing query.

You do this by placing the subquery in the FROM clause of the containing query as you would a table name. You may use subqueries in place of tables in this way as well in INSERT, UPDATE, and DELETE statements.

Subqueries so used can employ correlation variables, but only those defined within the subquery itself, not outer references. Outer references ("left-

correlated subqueries") are allowed only in the `FROM` clause of a `SELECT` statement. See [table\\_collection\\_expression](#) on page 7-576.

A **subquery** answers multiple-part questions. For example, to determine who works in Taylor's department, you can first use a subquery to determine the department in which Taylor works. You can then answer the original question with the parent `SELECT` statement.

A subquery can contain another subquery. You can nest up to 255 levels of subqueries.

If tables in a subquery have the same name as tables in the containing statement, you must prefix any reference to the column of the table from the containing statement with the table name or alias. To make your statements easier for you to read, always qualify the columns in a subquery with the name or alias of the table, view, or materialized view.

Oracle performs a **correlated subquery** when the subquery references a column from a table referred to in the parent statement. A correlated subquery is evaluated once for each row processed by the parent statement. The parent statement can be a `SELECT`, `UPDATE`, or `DELETE` statement. See the "[Correlated Subquery Examples](#)" on page 7-595.

A correlated subquery answers a multiple-part question whose answer depends on the value in each row processed by the parent statement. For example, you can use a correlated subquery to determine which employees earn more than the average salaries for their departments. In this case, the correlated subquery specifically computes the average salary for each department.

## Unnesting of Nested Subqueries

Subqueries are "nested" when they appear in the `WHERE` clause of the parent statement. When Oracle evaluates a statement with a nested subquery, it must evaluate the subquery portion multiple times and may overlook some efficient access paths or joins.

**Subquery unnesting** unnests and merges the body of the subquery into the body of the statement that contains it, allowing the optimizer to consider them together when evaluating access paths and joins. The optimizer can unnest most subqueries, with some exceptions. Those exceptions include subqueries that contain a `CONNECT BY` or `START WITH` clause, a `ROWNUM` pseudocolumn, one of the set operators, a nested aggregate function, or a correlated reference to a query block that is not the subquery's immediate outer query block.

Assuming no restrictions exist, the optimizer automatically unnests some nested subqueries:

- Uncorrelated `IN` subqueries
- `IN` and `EXISTS` correlated subqueries as long, as they do not contain aggregate functions or a `GROUP BY` clause.

You can enable **extended subquery unnesting** by instructing the optimizer to unnest additional types of subqueries:

- You can unnest an uncorrelated `NOT IN` subquery by specifying the `HASH_AJ` or `MERGE_AJ` hint in the subquery.
- You can unnest other subqueries by specifying the `UNNEST` hint in the subquery

For information on hints, see [Chapter 2, "Basic Elements of Oracle SQL"](#).

## Selecting from the DUAL Table

`DUAL` is a table automatically created by Oracle along with the data dictionary. `DUAL` is in the schema of the user `SYS`, but is accessible by the name `DUAL` to all users. It has one column, `DUMMY`, defined to be `VARCHAR2(1)`, and contains one row with a value 'X'. Selecting from the `DUAL` table is useful for computing a constant expression with the `SELECT` statement. Because `DUAL` has only one row, the constant is returned only once. Alternatively, you can select a constant, pseudocolumn, or expression from any table, but the value will be returned as many times as there are rows in the table. See ["SQL Functions"](#) on page 4-2 for many examples of selecting a constant value from `DUAL`.

## Distributed Queries

Oracle's distributed database management system architecture allows you to access data in remote databases using Net8 and an Oracle server. You can identify a remote table, view, or materialized view by appending *@dblink* to the end of its name. The *dblink* must be a complete or partial name for a database link to the database containing the remote table, view, or materialized view.

**See Also:** ["Referring to Objects in Remote Databases"](#) on page 2-79 for more information on referring to database links.

### Restrictions on Distributed Queries

Distributed queries are currently subject to the restriction that all tables locked by a `FOR UPDATE` clause and all tables with `LONG` columns selected by the query must be

located on the same database. For example, the following statement will raise an error:

```
SELECT emp_ny.*
       FROM emp_ny@ny, dept
       WHERE emp_ny.deptno = dept.deptno
       AND dept.dname = 'ACCOUNTING'
       FOR UPDATE OF emp_ny.sal;
```

The following statement fails because it selects `LONG_COLUMN`, a `LONG` value, from the `EMP_REVIEW` table on the `NY` database and locks the `EMP` table on the local database:

```
SELECT emp.empno, review.long_column, emp.sal
       FROM emp, emp_review@ny review
       WHERE emp.empno = emp_review.empno
       FOR UPDATE OF emp.sal;
```

In addition, Oracle currently does not support distributed queries that select user-defined types or object REFs on remote tables.



---

# About SQL Statements

*... he is eminently plain and direct ... both in his syntax and in his words ...*

Matthew Arnold, *On Translating Homer*

This chapter describes the various types of Oracle SQL statements, and provides guidelines for finding the right SQL statement for your task. Topics include:

- [Summary of SQL Statements](#)
- [Finding the Right SQL Statement](#)

## Summary of SQL Statements

The tables in the following sections provide a functional summary of SQL statements and are divided into these categories:

- Data Definition Language (DDL) Statements
- Data Manipulation Language (DML) Statements
- Transaction Control Statements
- Session Control Statements
- System Control Statements

### Data Definition Language (DDL) Statements

Data definition language (DDL) statements enable you to perform these tasks:

- Create, alter, and drop schema objects
- Grant and revoke privileges and roles
- Analyze information on a table, index, or cluster

- Establish auditing options
- Add comments to the data dictionary

The `CREATE`, `ALTER`, and `DROP` commands require exclusive access to the specified object. For example, an `ALTER TABLE` statement fails if another user has an open transaction on the specified table.

The `GRANT`, `REVOKE`, `ANALYZE`, `AUDIT`, and `COMMENT` commands do not require exclusive access to the specified object. For example, you can analyze a table while other users are updating the table.

Oracle implicitly commits the current transaction before and after every DDL statement.

Many DDL statements may cause Oracle to recompile or reauthorize schema objects. For information on how Oracle recompiles and reauthorizes schema objects and the circumstances under which a DDL statement would cause this, see *Oracle8i Concepts*.

DDL statements are supported by PL/SQL with the use of the `DBMS_SQL` package.

**See Also:** *Oracle8i Supplied PL/SQL Packages Reference*.

[Table 6–1](#) lists the DDL statements.



**Table 6–1 Data Definition Language Statements**

ALTER CLUSTER	CREATE DIMENSION	DROP DATABASE LINK
ALTER DATABASE	CREATE DIRECTORY	DROP DIMENSION
ALTER DIMENSION	CREATE FUNCTION	DROP DIRECTORY
ALTER FUNCTION	CREATE INDEX	DROP FUNCTION
ALTER INDEX	CREATE INDEXTYPE	DROP INDEX
ALTER MATERIALIZED VIEW / SNAPSHOT	CREATE LIBRARY	DROP INDEXTYPE
ALTER MATERIALIZED VIEW / SHAPSHOT LOG	CREATE MATERIALIZED VIEW / SHAPSHOT	DROP LIBRARY
ALTER PACKAGE	CREATE MATERIALIZED VIEW / SNAPSHOT LOG	DROP MATERIALIZED VIEW / SNAPSHOT
ALTER PROCEDURE	CREATE OPERATOR	DROP MATERIALIZED VIEW / SNAPSHOT LOG
ALTER PROFILE	CREATE PACKAGE	DROP OPERATOR
ALTER RESOURCE COST	CREATE PACKAGE BODY	DROP PACKAGE
ALTER ROLE	CREATE PROCEDURE	DROP PROCEDURE
ALTER ROLLBACK SEGMENT	CREATE PROFILE	DROP PROFILE
ALTER SEQUENCE	CREATE ROLE	DROP ROLE
ALTER SNAPSHOT	CREATE ROLLBACK SEGMENT	DROP ROLLBACK SEGMENT
ALTER SHAPSHOT LOG	CREATE SCHEMA	DROP SEQUENCE
ALTER TABLE	CREATE SEQUENCE	DROP SNAPSHOT
ALTER TABLESPACE	CREATE SHAPSHOT	DROP SNAPSHOT LOG
ALTER TRIGGER	CREATE SNAPSHOT LOG	DROP SYNONYM
ALTER TYPE	CREATE SYNONYM	DROP TABLE
ALTER USER	CREATE TABLE	DROP TABLESPACE
ALTER VIEW	CREATE TABLESPACE	DROP TRIGGER
ANALYZE	CREATE TEMPORARY TABLESPACE	DROP TYPE
ASSOCIATE STATISTICS	CREATE TRIGGER	DROP USER
AUDIT	CREATE TYPE	DROP VIEW
COMMENT	CREATE USER	GRANT
CREATE CLUSTER	CREATE VIEW	NOAUDIT
CREATE CONTEXT	DISASSOCIATE STATISTICS	RENAME
CREATE CONTROLFILE	DROP CLUSTER	REVOKE
CREATE DATABASE	DROP CONTEXT	TRUNCATE
CREATE DATABASE LINK		

## Data Manipulation Language (DML) Statements

Data manipulation language (DML) statements query and manipulate data in existing schema objects. These statements do not implicitly commit the current transaction.

**Table 6–2 Data Manipulation Language Statements**

<b>Statement</b>
CALL
DELETE
EXPLAIN PLAN
INSERT
LOCK TABLE
SELECT
UPDATE

The `CALL` and `EXPLAIN PLAN` statements are supported in PL/SQL only when executed dynamically. All other DML statements are fully supported in PL/SQL.

## Transaction Control Statements

Transaction control statements manage changes made by DML statements.

**Table 6–3 Transaction Control Statements**

<b>Statement</b>
COMMIT
ROLLBACK
SAVEPOINT
SET TRANSACTION

All transaction control statements except certain forms of the `COMMIT` and `ROLLBACK` commands are supported in PL/SQL. For information on the restrictions, see "[COMMIT](#)" on page 7-230 and "[ROLLBACK](#)" on page 7-564.

## Session Control Statements

Session control statements dynamically manage the properties of a user session. These statements do not implicitly commit the current transaction.

PL/SQL does not support session control statements.

**Table 6–4 Session Control Statements**

Statement
ALTER SESSION
SET ROLE

## System Control Statement

The single system control statement dynamically manages the properties of an Oracle instance. This statement does not implicitly commit the current transaction.

ALTER SYSTEM is not supported in PL/SQL.

**Table 6–5 System Control Statement**

Statement
ALTER SYSTEM

## Embedded SQL Statements

Embedded SQL statements place DDL, DML, and transaction control statements within a procedural language program. Embedded SQL is supported by the Oracle precompilers and is documented in the following books:

- *Pro\*COBOL Precompiler Programmer's Guide*
- *Pro\*C/C++ Precompiler Programmer's Guide*
- *SQL\*Module for Ada Programmer's Guide*

## Finding the Right SQL Statement

The particular SQL statement you use to accomplish a given database task is sometimes obvious and sometimes difficult to predict. For example, you create a table with the CREATE TABLE statement. However, you don't enable a constraint with the ENABLE CONSTRAINT statement, because such a statement doesn't exist. Rather, you modify the column options using the ALTER TABLE statement.

This section lists, by database object and task, the appropriate SQL statement to use to accomplish various database tasks. You can then refer to [Chapter 7, "SQL Statements"](#), for the syntax and semantics of each SQL statement.

---



---

**Note:** Your ability to use the SQL statements listed in this section depends on the version and edition of Oracle you are using, as well as the options you have installed. Be sure to read the detailed descriptions in [Chapter 7, "SQL Statements"](#), before using these statements.

---



---

Database Object / Task	Operation	SQL Statement
<b>application</b>	allowing to connect as a user	ALTER USER <i>proxy_clause</i>
<b>application server</b>	allowing to connect as a user	ALTER USER <i>proxy_clause</i>
<b>auditing</b>	of database events	CREATE TRIGGER
<b>call</b>	limit CPU time for	CPU_PER_CALL <i>parameter</i>
	limit data blocks read	LOGICAL_READS_PER_CALL <i>parameter</i>
<b>checkpoint</b>	perform explicitly	ALTER SYSTEM CHECKPOINT
<b>clone database</b>	mount	ALTER DATABASE MOUNT
<b>cluster</b>	cluster key, change columns of	prohibited
	extent, allocate for	ALTER CLUSTER <i>allocate_extent_clause</i>
	migrated or chained rows, identify	ANALYZE
	parallelism of, change	ALTER CLUSTER <i>parallel_clause</i>
	rename	prohibited
	storage characteristics of, change	ALTER CLUSTER <i>physical_attributes_clause</i>
	tablespace of, change	prohibited
	unused space in, release	ALTER CLUSTER <i>deallocate_unused_clause</i>
<b>column</b>	add to a table or modify	ALTER TABLE <i>add_column_options</i> , <i>modify_column_options</i>
	define	CREATE TABLE

Database Object / Task	Operation	SQL Statement
	drop from a table	ALTER TABLE <i>drop_column_clause</i>
	generate derived values automatically	CREATE TRIGGER
	organization of, define	CREATE TABLE
<b>commit operation</b>	prevent procedure or function from issuing	ALTER SESSION
<b>compilation</b>	avoid run-time of	ALTER FUNCTION ... COMPILE
<b>constraint</b>	add to a table or modify	ALTER TABLE <i>add_column_options</i> , <i>modify_column_options</i>
	business, enforce	CREATE TRIGGER
	enable, disable, or drop	ALTER TABLE <i>enable_disable_clause</i> , <i>drop_constraint_clause</i>
	specify	CREATE TABLE
<b>control file</b>	back up	ALTER DATABASE <i>controlfile_clauses</i>
	standby, create	ALTER DATABASE CREATE STANDBY CONTROLFILE
<b>currency symbol</b>	reset for session	ALTER SESSION SET NLS_CURRENCY
<b>data</b>	frequently used, caching	ALTER TABLE <i>cache_clause</i>
	specify as temporary or permanent	CREATE TABLE
<b>data dictionary</b>	convert from Oracle7 to Oracle8i	ALTER DATABASE CONVERT
<b>data independence</b>	provide	CREATE SYNONYM
<b>database</b>	character set of, change	ALTER DATABASE CHARACTER SET
	create script for	ALTER DATABASE <i>controlfile_clauses</i>
	database character set for, specify	CREATE DATABASE
	datafiles for, specify	CREATE DATABASE
	datafiles of, modify	ALTER DATABASE
	datafiles, establish number of	CREATE DATABASE
	downgrade to an earlier release	ALTER DATABASE RESET COMPATIBILITY
	global name of, change	ALTER DATABASE RENAME GLOBAL_NAME

<b>Database Object / Task</b>	<b>Operation</b>	<b>SQL Statement</b>
	global name resolution, enable for the session	ALTER SESSION SET GLOBAL_NAMES
	instances, establish number of	CREATE DATABASE
	media recovery, design	ALTER DATABASE <i>general_recovery_clause</i>
	media recovery, perform ongoing	ALTER DATABASE <i>managed_recovery_clause</i>
	mount	ALTER DATABASE MOUNT
	move a subset to a different Oracle database	ALTER TABLE <i>exchange_partition_clause</i>
	national character set for, specify	CREATE DATABASE
	national character set of, change	ALTER DATABASE CHARACTER SET
	open	ALTER DATABASE OPEN
	parallelize recovery of	ALTER DATABASE <i>parallel_clause</i>
	place in read-only mode	ALTER DATABASE OPEN
	place in read-write mode	ALTER DATABASE OPEN
	place in sustained standby recovery mode	ALTER DATABASE <i>general_recovery_clause</i>
	prepare to re-create	ALTER DATABASE <i>controlfile_clauses</i>
	recover	ALTER DATABASE <i>recover_clauses</i>
	redo log file groups, establish number of	CREATE DATABASE
	redo log files for, specify	CREATE DATABASE
	redo log files of, create or modify	ALTER DATABASE
	redo log files, establish number of	CREATE DATABASE
	redo log, choose mode for	CREATE DATABASE
	upgrade to Oracle8i	ALTER DATABASE
<b>database character set</b>	specify for a database	CREATE DATABASE
<b>database events</b>	transparent logging of	CREATE TRIGGER
<b>database link</b>	close	ALTER SESSION

Database Object / Task	Operation	SQL Statement
<b>database security</b>	enforce authorizations	CREATE TRIGGER
<b>datafile</b>	automatic extension of, allow	ALTER DATABASE DATAFILE <i>autoextend_clause</i>
	create	ALTER DATABASE CREATE DATAFILE
	put online	ALTER DATABASE DATAFILE ONLINE
	reconstruct damaged	ALTER DATABASE <i>general_recovery_clause</i>
	reconstruct lost or damaged	ALTER DATABASE CREATE DATAFILE
	recover specified	ALTER DATABASE <i>general_recovery_clause</i>
	replace an old, for recovery	ALTER DATABASE CREATE DATAFILE
	resize	ALTER DATABASE DATAFILE RESIZE
	take offline	ALTER DATABASE DATAFILE ONLINE/ OFFLINE
	begin or end backup of	ALTER TABLESPACE ... BACKUP
	number of, establish for a database	CREATE DATABASE
	online, update instance information on	ALTER SYSTEM <i>check_datafiles_clause</i>
	specify for a database	CREATE DATABASE
<b>dates</b>	format of	See Table 2-9, " <a href="#">Date Format Elements</a> " on page 2-45.
<b>decimal character</b>	reset for session	ALTER SESSION SET NLS_NUMERIC_CHARACTERS
<b>dimension</b>	add a level, hierarchy, or attribute to	ALTER DIMENSION ... ADD
	change the relationships of	ALTER DIMENSION
	drop a level, hierarchy, or attribute from	ALTER DIMENSION ... DROP
	explicitly compile	ALTER DIMENSION ... COMPILE
<b>dispatcher processes</b>	multi-threaded server, manage	MTS_ parameters of ALTER SYSTEM

Database Object / Task	Operation	SQL Statement
<b>domain index</b>	alter	ALTER INDEX ... PARAMETERS
	rebuild	ALTER INDEX <i>rebuild_clause</i>
<b>dump file</b>	limit the size of	ALTER SESSION SET MAX_DUMP_FILE_SIZE
<b>error messages</b>	language in which displayed, change	ALTER SESSION SET NLS_LANGUAGE
<b>function</b>	allow to or prevent from committing a transaction	ALTER SESSION
	declaration of, change	CREATE OR REPLACE FUNCTION
	definition of, change	CREATE OR REPLACE FUNCTION
	recompile explicitly	ALTER FUNCTION
<b>function-based index</b>	disable	ALTER INDEX ... [ <i>rebuild_clause</i> ] DISABLE
	disabled, re-enable	ALTER INDEX ... [ <i>rebuild_clause</i> ] ENABLE
<b>global names</b>	enforce resolution of	GLOBAL_NAMES parameter of ALTER SYSTEM
<b>hash join operations</b>	data blocks for, allocate	ALTER SESSION SET HASH_MULTIBLOCK_IO_COUNT
	in queries, enable or disable	ALTER SESSION SET HASH_JOIN_ENABLED ...
	memory for, allocate	ALTER SESSION SET HASH_AREA_SIZE
<b>index</b>	allow DML operations during rebuilding of	ALTER INDEX <i>rebuild_clause</i>
	based on a function; see "function-based index"	CREATE INDEX ... <i>column_expression</i>
	based on an indextype; see "domain index"	CREATE INDEX <i>domain_index_clause</i>
	collect statistics during rebuilding of	ALTER INDEX <i>rebuild_clause</i>
	default attribute values of, change	ALTER INDEX <i>partitioning_clauses</i>
	degree of parallelism for, change	ALTER INDEX <i>parallel_clause</i>
	direct-load INSERT operations, write to a log	ALTER INDEX <i>physical_attributes_clause</i>



Database Object / Task	Operation	SQL Statement
	extent for, allocate new	ALTER INDEX <i>allocate_extent_clause</i>
	key compression, enable	ALTER INDEX <i>rebuild_clause</i>
	key values, eliminate repetition of	ALTER INDEX <i>rebuild_clause</i>
	merge block contents of	ALTER INDEX <i>rebuild_clause</i>
	physical attributes of a partition of, change	ALTER INDEX <i>physical_attributes_clause</i>
	physical attributes of a subpartition of, change the	ALTER INDEX <i>physical_attributes_clause</i>
	physical attributes of, change	ALTER INDEX <i>physical_attributes_clause</i>
	re-create	ALTER INDEX <i>rebuild_clause</i>
	rebuild operations, write to a log	ALTER INDEX <i>rebuild_clause</i>
	SQL*Loader operations against, write to a log	ALTER INDEX <i>physical_attributes_clause</i>
	store bytes in reverse order	ALTER INDEX <i>rebuild_clause</i>
	tablespace for, specify	ALTER INDEX <i>rebuild_clause</i>
	tell Oracle not to use	ALTER INDEX ... [ <i>rebuild_clause</i> ] UNUSABLE
	unused space, release	ALTER INDEX <i>deallocate_unused_clause</i>
	rename	ALTER INDEX <i>rebuild_clause</i>
<b>index partition</b>	create-time attributes, change	ALTER INDEX <i>rebuild_clause</i>
	log direct-load INSERT operations	ALTER INDEX <i>physical_attributes_clause</i>
	log SQL*Loader operations against	ALTER INDEX <i>physical_attributes_clause</i>
	move to a different tablespace	ALTER INDEX <i>rebuild_clause</i>
	physical attributes of, change	ALTER INDEX <i>physical_attributes_clause</i>
	physical, logging, or storage characteristics of, change	ALTER INDEX <i>partitioning_clauses</i>
	re-create	ALTER INDEX <i>rebuild_clause</i>

Database Object / Task	Operation	SQL Statement
	remove from the database	ALTER INDEX <i>partitioning_clauses</i>
	specify a tablespace for	ALTER INDEX <i>rebuild_clause</i>
	split into two partitions	ALTER INDEX <i>partitioning_clauses</i>
	tell Oracle not to use	ALTER INDEX ... UNUSABLE
<b>index subpartition</b>	change a create-time attributes, change	ALTER INDEX <i>rebuild_clause</i>
	log direct-load INSERT operations	ALTER INDEX <i>physical_attributes_clause</i>
	log SQL*Loader operations against	ALTER INDEX <i>physical_attributes_clause</i>
	move to a different tablespace	ALTER INDEX <i>rebuild_clause</i>
	physical attributes, change	ALTER INDEX <i>physical_attributes_clause</i>
	physical, logging, or storage characteristics, change	ALTER INDEX <i>partitioning_clauses</i>
	re-create	ALTER INDEX <i>rebuild_clause</i>
	tablespace for, specify	ALTER INDEX <i>rebuild_clause</i>
	tell Oracle not to use	ALTER INDEX ... UNUSABLE
<b>index-organized table</b>	characteristics, change	ALTER TABLE
<b>indexes</b>	on a cluster	CREATE INDEX
	on a nested table storage table	CREATE INDEX
	on a partitioned table	CREATE INDEX
	on an index-organized table	CREATE INDEX
	on columns of a table	CREATE INDEX
	on scalar typed object attributes	CREATE INDEX
<b>instance</b>	dynamically modify	ALTER SYSTEM
	make an index extent available to	ALTER INDEX <i>allocate_extent_clause</i>
	switch to a different	ALTER SESSION SET INSTANCE
<b>instance recovery</b>	continue after interruption	ALTER DATABASE <i>general_recovery_clause</i>

Database Object / Task	Operation	SQL Statement
<b>instances</b>	number of, establish for a database	CREATE DATABASE
<b>Java class</b>	force resolution of	ALTER JAVA
<b>Java resource</b>	force compilation of	ALTER JAVA
<b>Java source</b>	force compilation of	ALTER JAVA
<b>licensing</b>	changing limits or thresholds	LICENSE_ parameters of ALTER SYSTEM
<b>LOB columns</b>	add to a table or modify	ALTER TABLE <i>add_column_options</i> , <i>modify_column_options</i> , <i>LOB_storage_clause</i>
<b>location transparency</b>	provide	CREATE SYNONYM
<b>materialized view</b>	automatic refresh, change the mode or timing of	ALTER MATERIALIZED VIEW <i>refresh_clause</i>
	change from rowid-based to primary-key-based	ALTER MATERIALIZED VIEW ALTER MATERIALIZED VIEW LOG
	degree of parallelism, specify or change	ALTER MATERIALIZED VIEW <i>parallel_clause</i>
	divide into partitions	ALTER MATERIALIZED VIEW <i>partitioning_clauses</i>
	LOB storage characteristics, change	ALTER MATERIALIZED VIEW <i>modify_LOB_storage_clause</i>
	LOB storage characteristics, specify	ALTER MATERIALIZED VIEW <i>LOB_storage_clause</i>
	log changes to	ALTER MATERIALIZED VIEW ... LOGGING
	make eligible for query rewrite	ALTER MATERIALIZED VIEW ... QUERY REWRITE ALTER SESSION SET QUERY_REWRITE_ENABLED
	make frequently accessed data accessible	ALTER MATERIALIZED VIEW ... CACHE
	revalidate	ALTER MATERIALIZED VIEW ... COMPILE
	storage characteristics, change	ALTER MATERIALIZED VIEW <i>physical_attributes_clause</i>

Database Object / Task	Operation	SQL Statement
<b>materialized view log</b>	automatic refresh, change the mode and timing of	ALTER MATERIALIZED VIEW LOG
	change from rowid-based to primary-key-based	ALTER MATERIALIZED VIEW LOG
	divide into partitions	ALTER MATERIALIZED VIEW LOG <i>partitioning_clauses</i>
	physical and storage characteristics, change	ALTER MATERIALIZED VIEW LOG ... <i>physical_attributes_clause</i>
	save both old and new values	ALTER MATERIALIZED VIEW LOG ...NEW VALUES
	store primary key of changed rows	ALTER MATERIALIZED VIEW LOG ... ADD
	store rowid of changed rows	ALTER MATERIALIZED VIEW LOG ... ADD
<b>media recovery</b>	avoid on startup	ALTER DATABASE DATAFILE END BACKUP
	from specified redo log file	ALTER DATABASE <i>general_recovery_clause</i>
	prepare for	ALTER DATABASE ARCHIVELOG
<b>national character set</b>	specify for a database	CREATE DATABASE
<b>national language support</b>	change settings for the session	ALTER SESSION SET NLS_ parameters
<b>nested table</b>	update in a view	create an INSTEAD OF trigger
<b>nested table columns</b>	indexing	CREATE INDEX
<b>numbers</b>	format	See Table 2-7, " <a href="#">Number Format Elements</a> " on page 2-41.
<b>object references. See REFs</b>		
<b>online redo log</b>	reinitialize	ALTER DATABASE CLEAR LOGFILE
<b>outline</b>	assign to a different category	ALTER OUTLINE ... CHANGE CATEGORY TO
	recompile	ALTER OUTLINE ... REBUILD
	rename	ALTER OUTLINE ... RENAME
	automatically create and store	ALTER SESSION SET CREATE_STORED_OUTLINES

Database Object / Task	Operation	SQL Statement
	use to generate execution plans	ALTER SESSION SET USE_STORED_OUTLINES
<b>package</b>	avoid run-time compilation	ALTER PACKAGE
	compile explicitly	ALTER PACKAGE
<b>package body</b>	avoid run-time compilation	ALTER PACKAGE
	recompile explicitly	ALTER PACKAGE
<b>parallelism</b>	specify for a table	CREATE TABLE
	specify for DML on a table	CREATE TABLE
<b>parameter, initialization</b>	change the setting for the current session	ALTER SESSION <i>set_clause</i>
<b>parameter, session</b>	set or change the setting of	ALTER SESSION <i>set_clause</i>
<b>partition</b>	add to a table or modify	ALTER TABLE
	default attributes, change	ALTER TABLE <i>modify_default_attributes_clause</i>
	logging characteristics, change	ALTER TABLE <i>logging_clause</i>
	merge with another partition	ALTER TABLE <i>merge_partitions_clause</i>
	point to data in a nonpartitioned table	ALTER TABLE <i>exchange_partition_clause</i>
	real attributes, change	ALTER TABLE <i>modify_partition_clause</i>
<b>password</b>	complexity of, guarantee	PASSWORD_VERIFY_FUNCTION parameter
	make unavailable	PASSWORD_REUSE_TIME parameter
	number of days account will be locked after failed login attempts, specify	PASSWORD_LOCK_TIME parameter
	number of days before reuse, limit	PASSWORD_REUSE_TIME parameter
	number of days in grace period, specify	PASSWORD_GRACE_TIME parameter
	number of days usable, limit	PASSWORD_LIFE_TIME parameter
	number of times reused, limit	PASSWORD_REUSE_MAX parameter
	special characters in, allow	PASSWORD_VERIFY_FUNCTION parameter

Database Object / Task	Operation	SQL Statement
<b>performance</b>	optimize for index access path	ALTER SESSION SET OPTIMIZER_INDEX_COST_ADJ
	optimize for nested loop joins	ALTER SESSION SET OPTIMIZER_INDEX_CACHING
	specify the optimizer approach for the session	ALTER SESSION SET OPTIMIZER_MODE
<b>procedure</b>	allow to or prevent from committing a transaction	ALTER SESSION
	avoid run-time compilation	ALTER PROCEDURE
	recompile explicitly	ALTER PROCEDURE
<b>profile</b>	resource limit, add to	ALTER PROFILE
	resource limit, change	ALTER PROFILE
	resource limit, drop from	ALTER PROFILE
<b>recovery</b>	distributed, enable or disable	ALTER SYSTEM <i>distributed_recovery_clause</i>
<b>recovery data</b>	discard	ALTER DATABASE RESETLOGS
<b>redo log</b>	remove changes from	ALTER DATABASE OPEN RESETLOGS
	reset sequence of	ALTER DATABASE OPEN RESETLOGS
	specify mode of	CREATE DATABASE
<b>redo log file</b>	add	ALTER DATABASE ADD LOGFILE MEMBER
	automatically generates names for	ALTER DATABASE <i>general_recovery_clause</i>
	clear	ALTER DATABASE CLEAR LOGFILE
	drop	ALTER DATABASE DROP LOGFILE
	enable or disable thread	ALTER DATABASE ENABLE THREAD
	rename	ALTER DATABASE RENAME FILE
	number of, establish for a database	CREATE DATABASE
	archive manually or automatically	ALTER SYSETM <i>archive_log_clause</i>
	number of, establish for a database	CREATE DATABASE

Database Object / Task	Operation	SQL Statement
	specify a path for	ALTER SESSION SET LOG_ARCHIVE_DEST_n
	switch manually	ALTER SYSTEM <i>switch_logfile_clause</i>
<b>REFS</b>	validate and update	ANALYZE
<b>role</b>	change authorization required	ALTER ROLE
<b>rollback segment</b>	bring online	ALTER ROLLBACK SEGMENT
	reduce in size	ALTER ROLLBACK SEGMENT
	storage characteristics, change	ALTER ROLLBACK SEGMENT
	take offline	ALTER ROLLBACK SEGMENT
<b>rowid</b>	examine	query the ROWID pseudocolumn
	extended, interpreting contents	DBMS_ROWID package; see <i>Oracle8i Supplied PL/SQL Packages Reference</i>
<b>schema</b>	change during the session	ALTER SESSION SET CURRENT_SCHEMA
<b>schema object</b>	reference without referencing its location	CREATE SYNONYM
	reference without referencing its owner	CREATE SYNONYM
	specify another name for	CREATE SYNONYM
	validate structure of	ANALYZE
<b>sequence</b>	cached sequence values, change number of	ALTER SEQUENCE <i>cache_clause</i>
	consecutive order of values, guarantee	CREATE SEQUENCE ... ORDER ALTER SEQUENCE ... ORDER
	create	CREATE SEQUENCE
	determine current value of	See " <a href="#">CURRVAL and NEXTVAL</a> " on page 2-56.
	increment value, set	CREATE SEQUENCE ... INCREMENT BY ALTER SEQUENCE ... INCREMENT BY
	maximum or minimum value, eliminate	ALTER SEQUENCE
	minimum or maximum value, set	CREATE SEQUENCE ALTER SEQUENCE

Database Object / Task	Operation	SQL Statement
	preallocate values for faster access	CREATE SEQUENCE ALTER SEQUENCE
	restart after a predefined limit	CREATE SEQUENCE ... CYCLE ALTER SEQUENCE ... CYCLE
	starting value, set	CREATE SEQUENCE
<b>server processes</b>	multi-threaded server, manage	MTS_ parameters of ALTER SYSTEM
<b>session</b>	CPU time for, limit	CPU_PER_SESSION parameter
	data blocks read, limit	LOGICAL_READS_PER_SESSION parameter
	enable or disable parallel transactions in	ALTER SESSION
	inactive period duration, limit	IDLE_TIME parameter
	private SGA space for, limit	PRIVATE_SGA parameter
	resource costs allowed, change	ALTER RESOURCE COST
	restrict to privileged users	ALTER SYSTEM <i>restricted_session_clause</i>
	terminate	ALTER SYSTEM <i>kill_session_clause</i>
	total elapsed time, limit	CONNECT_TIME parameter
	total resources for, limit	COMPOSITE_LIMIT parameter
<b>SGA</b>	flush data from shared pool	ALTER SYSTEM <i>flush_shared_pool_clause</i>
<b>shared pool</b>	flush	ALTER SYSTEM <i>flush_shared_pool_clause</i>
	<b>snapshot.</b> See "materialized view".	
<b>sort operations</b>	linguistic sequence, change	ALTER SESSION SET NLS_SORT
<b>standby database</b>	activate	ALTER DATABASE ACTIVATE STANDBY DATABASE
	recover	ALTER DATABASE <i>recover_clauses</i>
<b>statistics</b>	on a schema object, collect	ANALYZE
	on a schema object, delete	ANALYZE
	on scalar object attributes, collect	ANALYZE
<b>subpartition</b>	add to a table or modify	ALTER TABLE



Database Object / Task	Operation	SQL Statement
	default attributes, change	ALTER TABLE <i>modify_default_attributes_clause</i> , <i>modify_partition_clause</i>
	logging characteristics, change	ALTER TABLE <i>logging_clause</i>
	real attributes, change	ALTER TABLE <i>modify_subpartition_clause</i>
<b>system resources</b>	enable or disable	RESOURCE_LIMITS parameter of ALTER SYSTEM
<b>table</b>	allocate space for	ALTER TABLE <i>allocate_extent_clause</i>
	characteristics, change	ALTER TABLE <i>physical_attributes_clause</i> , <i>modify_storage_clauses</i>
	column, drop from table	ALTER TABLE <i>drop_column_clause</i>
	degree of parallelism, change	ALTER TABLE <i>parallel_clause</i>
	logging characteristics, change	ALTER TABLE <i>logging_clause</i>
	make read-only, read-write	ALTER TABLE
	migrated or chained rows, identify	ANALYZE
	organization, define	CREATE TABLE
	partition, point to the contents of another table	ALTER TABLE <i>exchange_partition_clause</i>
	partitioning, specify	CREATE TABLE
	rename	ALTER TABLE
	unused space of, release	ALTER TABLE <i>deallocate_unused_clause</i>
	heap or index organized	CREATE TABLE
	include in a cluster	CREATE TABLE
	replicate asynchronous, maintain	CREATE TRIGGER
	storage characteristics of, set	CREATE TABLE
<b>tablespace</b>	allow or disallow writing to	ALTER TABLESPACE READ WRITE/ONLY
	datafiles, add or rename	ALTER TABLESPACE <i>datafile/</i> <i>tempfile_clauses</i>

Database Object / Task	Operation	SQL Statement
	logging characteristics, change	ALTER TABLESPACE
	minimum extent length, change	ALTER TABLESPACE
	reconstruct damaged	ALTER DATABASE <i>general_recovery_clause</i>
	reconstruct lost or damaged	ALTER DATABASE CREATE DATAFILE
	recover specified	ALTER DATABASE <i>general_recovery_clause</i>
	specifying for a table	CREATE TABLE
	storage characteristics, change	ALTER TABLESPACE
	take online or offline	ALTER TABLESPACE
	user quota on, change	ALTER USER
	assign to a user	CREATE USER
	space quota for a user, allocate	CREATE USER
<b>tempfile</b>	allow for automatic extension of	ALTER DATABASE TEMPFILE
	resize	ALTER DATABASE TEMPFILE
<b>transaction</b>	distributed, force commit of	ALTER SESSION
	distributed, force rollback of	ALTER SESSION
<b>trigger</b>	enable or disable	ALTER TABLE
<b>user</b>	authentication, change	ALTER USER
	database resources limits, change	ALTER USER <i>profile_clause</i>
	default roles, change	ALTER USER
	failed attempts to log in, limit	FAILED_LOGIN_ATTEMPTS parameter
	number of sessions, limit	SESSIONS_PER_USER parameter
	password, change	ALTER USER
	resource limits, set	CREATE USER

---

<b>Database Object / Task</b>	<b>Operation</b>	<b>SQL Statement</b>
	restrict access to Oracle	ALTER SYSTEM <i>restricted_session_clause</i>
	tablespace quota, allocate	CREATE USER
	tablespaces, assign	CREATE USER

---



---

---

# SQL Statements

*A whole is that which has a beginning, a middle, and an end.*

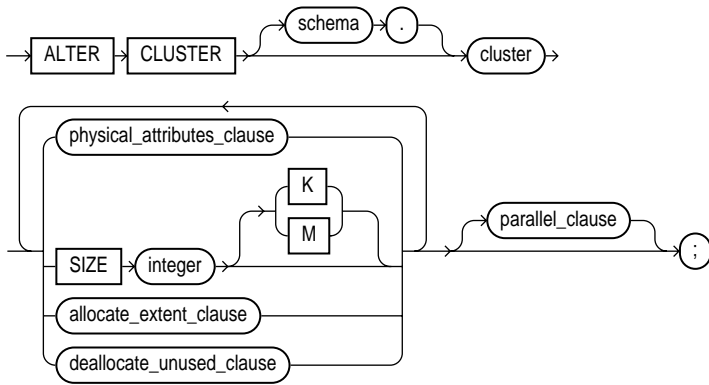
Aristotle, *Poetics*

This chapter describes, in alphabetical order, Oracle SQL statements and major clauses. The description of each statement or clause contains the following sections:

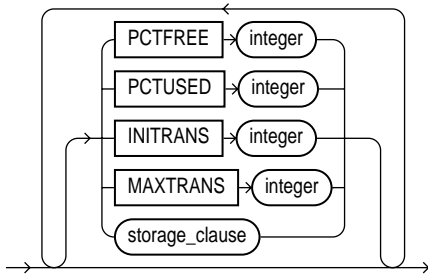
<b>Syntax</b>	shows the keywords and parameters that make up the statement.  <b>Caution: Not all keywords and parameters are valid in all circumstances. Be sure to refer to the "Keywords and Parameters" section of each statement and clause to learn about any restrictions on the syntax.</b>
<b>Purpose</b>	describes the basic uses of the statement.
<b>Prerequisites</b>	lists privileges you must have and steps that you must take before using the statement. In addition to the prerequisites listed, most statements also require that the database be opened by your instance, unless otherwise noted.
<b>Keywords and Parameters</b>	describes the purpose of each keyword and parameter. (The conventions for keywords and parameters used in this chapter are explained in the Preface of this reference.) Restrictions and usage notes also appear in this section.
<b>Examples</b>	shows how to use various clauses and parameters of the statement.

## ALTER CLUSTER

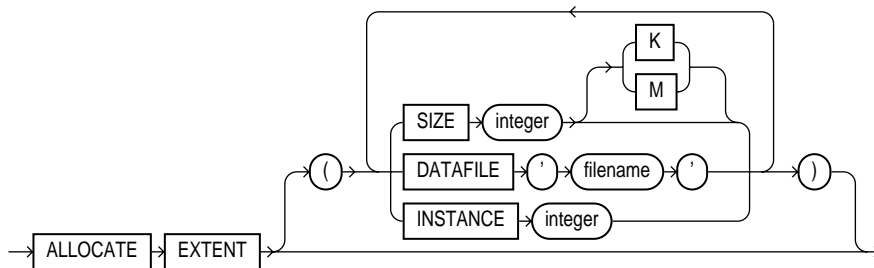


### Syntax



### physical\_attributes\_clause::=



**storage\_clause:** See "storage\_clause" on page 7-605.

**allocate\_extent\_clause::=****deallocate\_unused\_clause::=****parallel\_clause::=****Purpose**

To redefine storage and parallelism characteristics of a cluster.

For information on creating a cluster, see "[CREATE CLUSTER](#)" on page 7-254.

To remove tables from a cluster, see "[DROP CLUSTER](#)" on page 7-472 and "[DROP TABLE](#)" on page 7-501.

---



---

**Note:** You cannot use this statement to change the number or the name of columns in the cluster key, and you cannot change the tablespace in which the cluster is stored.

---



---

## Prerequisites

The cluster must be in your own schema or you must have ALTER ANY CLUSTER system privilege.

## Keywords and Parameters

---

<i>schema</i>	is the schema containing the cluster. If you omit <i>schema</i> , Oracle assumes the cluster is in your own schema.
<i>cluster</i>	is the name of the cluster to be altered.
<i>physical_attributes_clause</i>	changes the values of the PCTUSED, PCTFREE, INITTRANS, and MAXTRANS parameters of the cluster. For a description of these parameters, see "CREATE CLUSTER" on page 7-254.
<i>storage_clause</i>	changes the storage characteristics for the cluster. See the " <a href="#">storage_clause</a> " on page 7-605. <b>Restriction:</b> You cannot change the values of the storage parameters INITIAL and MINEXTENTS for a cluster.
SIZE <i>integer</i>	determines how many cluster keys will be stored in data blocks allocated to the cluster. For a description of the SIZE parameter, see "CREATE CLUSTER" on page 7-254. <b>Restriction:</b> You can change the SIZE parameter only for an indexed cluster, not for a hash cluster.
<i>allocate_extent_clause</i>	explicitly allocates a new extent for the cluster. <b>Restriction:</b> You can allocate a new extent only for an indexed cluster, not for a hash cluster.
SIZE	specifies the size of the extent in bytes. Use K or M to specify the extent size in kilobytes or megabytes.  When you explicitly allocate an extent with this clause, Oracle does not evaluate the cluster's storage parameters and determine a new size for the next extent to be allocated (as it does when you create a table). Therefore, specify SIZE if you do not want Oracle to use a default value.
DATAFILE	specifies one of the datafiles in the cluster's tablespace to contain the new extent. If you omit this parameter, Oracle chooses the datafile.
INSTANCE	makes the new extent available to the specified instance. An instance is identified by the value of its initialization parameter INSTANCE_NUMBER. If you omit INSTANCE, the extent is available to all instances. <b>Use this parameter only if you are using Oracle with the Parallel Server option in parallel mode.</b>



---

<i>deallocate_</i> <i>unused_clause</i>	explicitly deallocates unused space at the end of the cluster and makes the freed space available for other segments. Only unused space above the high water mark can be freed.
KEEP	specifies the number of bytes above the high water mark that the cluster will have after deallocation. If the number of remaining extents is less than MINEXTENTS, then MINEXTENTS is set to the current number of extents. If the initial extent becomes smaller than INITIAL, then INITIAL is set to the value of the current initial extent. If you omit KEEP, all unused space is freed.
	For a more complete description of this clause, see " <a href="#">ALTER TABLE</a> " on page 7-123.
<i>parallel_clause</i>	changes the default degree of parallelism for queries and DML on the cluster. <b>See Also:</b> The <a href="#">Notes</a> to the <i>parallel_clause</i> of " <a href="#">CREATE TABLE</a> " on page 4-381.
NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the PARALLEL_THREADS_PER_CPU initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution processes. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .
	<b>Restriction:</b> If the tables in <i>cluster</i> contain any columns of LOB or user-defined object type, this statement as well as subsequent INSERT, UPDATE, or DELETE operations on <i>cluster</i> are executed serially without notification.

---

## Examples

The following statement alters the CUSTOMER cluster in the schema SCOTT:

```
ALTER CLUSTER scott.customer
  SIZE 512
  STORAGE (MAXEXTENTS 25);
```

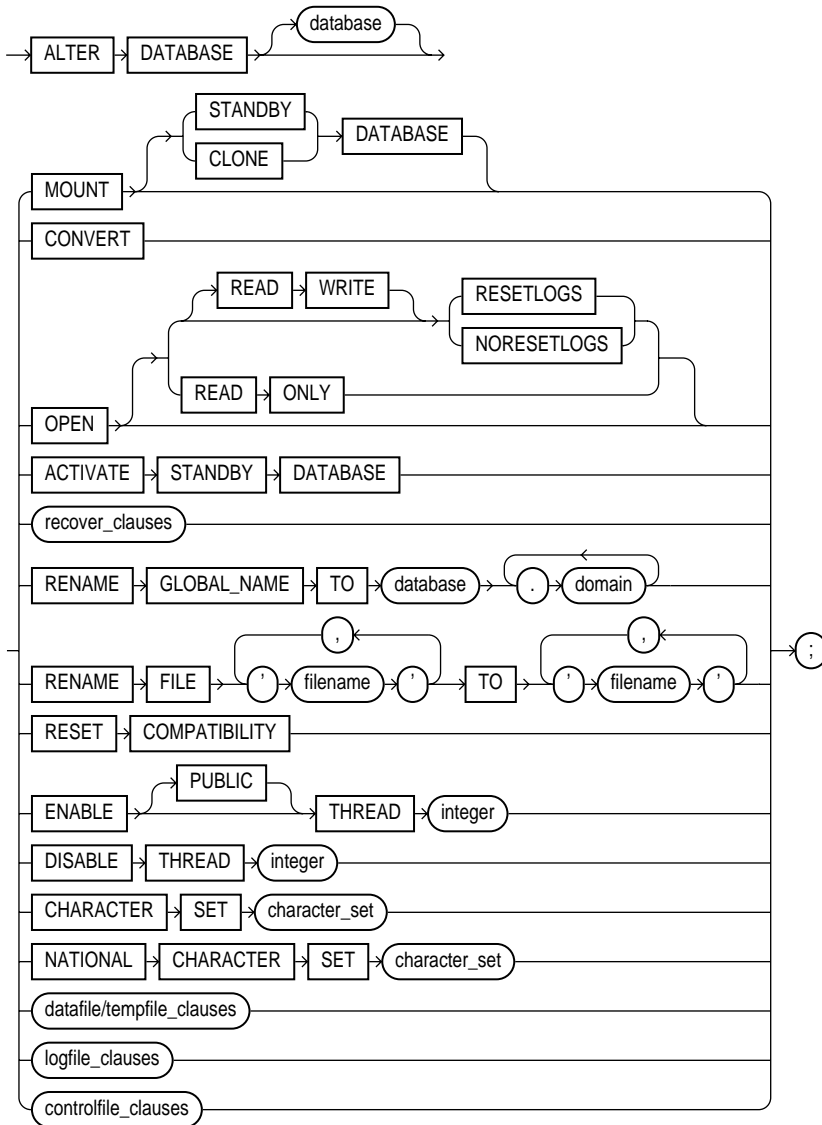
Oracle allocates 512 bytes for each cluster key value. Assuming a data block size of 2 kilobytes, future data blocks within this cluster contain 4 cluster keys per data block, or 2 kilobytes divided by 512 bytes. The cluster can have a maximum of 25 extents.

The following statement deallocates unused space from the CUSTOMER cluster, keeping 30 kilobytes of unused space for future use:


```
ALTER CLUSTER scott.customer
  DEALLOCATE UNUSED KEEP 30 K;
```

# ALTER DATABASE

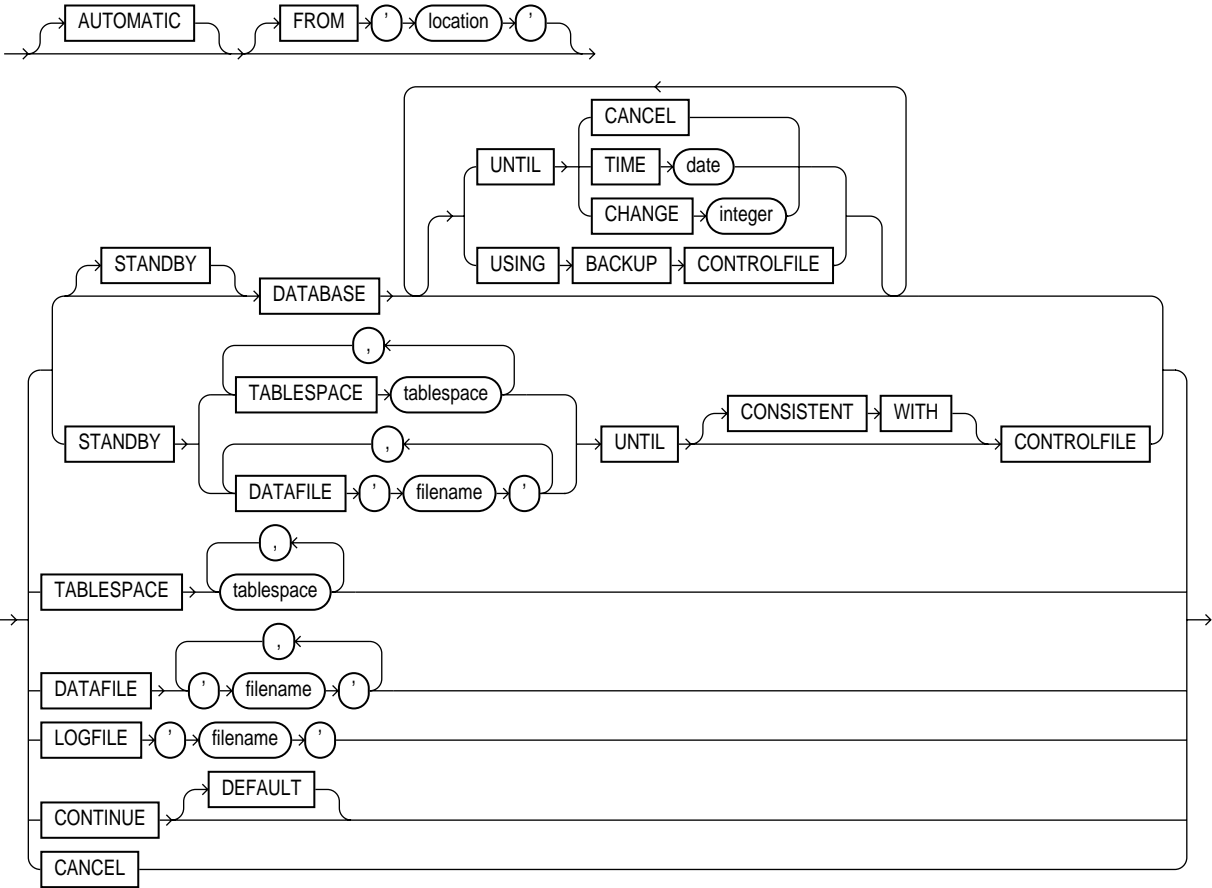
## Syntax




**recover\_clauses::=**



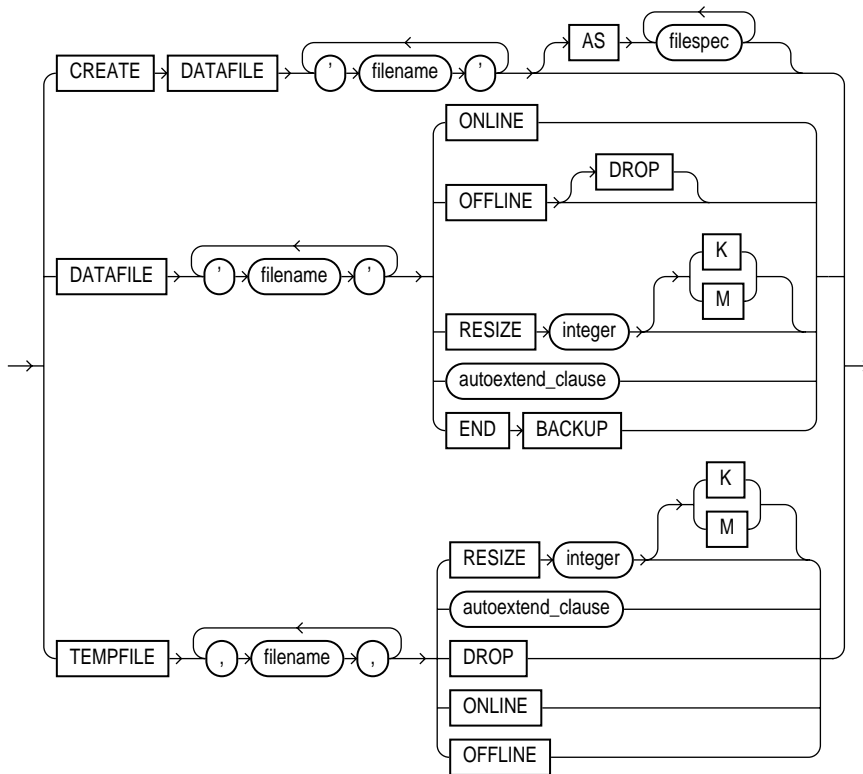
**general\_recovery\_clause::=**



**managed\_recovery\_clause::=**

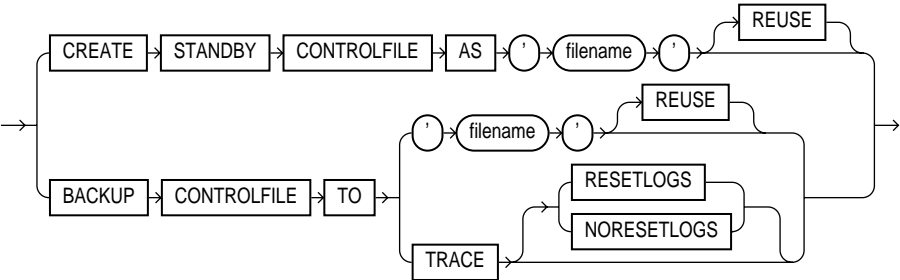


**datafile/tempfile\_clauses::=**

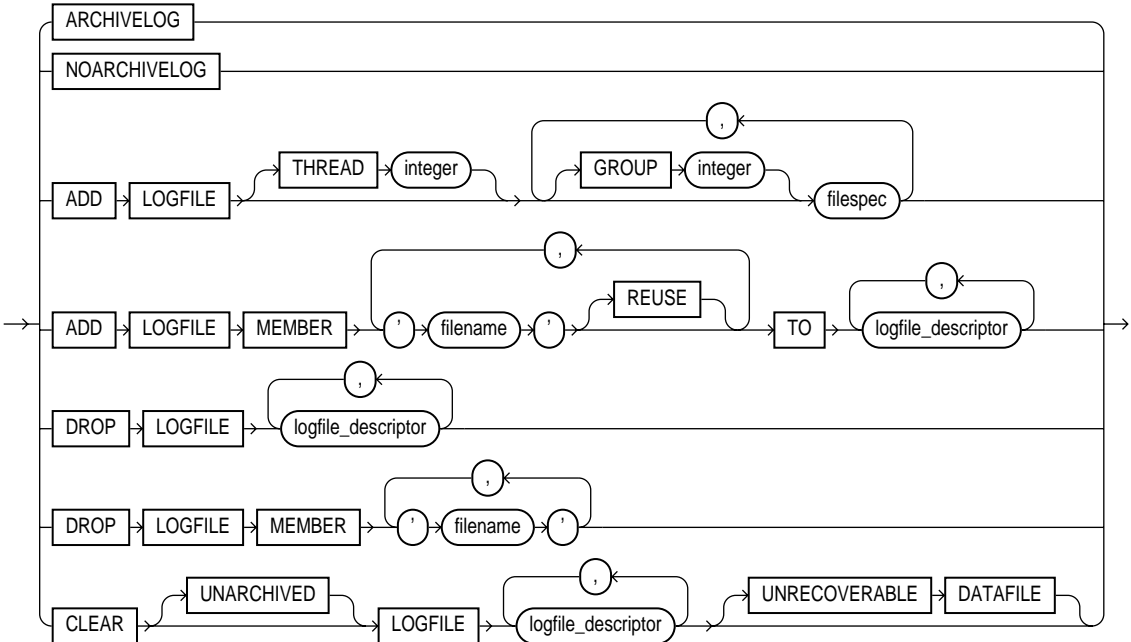


**filespec:** See "filespec" on page 7-516.


**controlfile\_clauses::=**




**logfile\_clauses::=**




**logfile\_descriptor::=**




**autoextend\_clause::=**



**maxsize\_clause::=**



**parallel\_clause::=**



**Purpose**

To modify, maintain, or recover an existing database.

**See Also:**

- *Oracle8i Administrator's Guide* for more information on using the ALTER DATABASE statement for database maintenance.
- *Oracle8i Administrator's Guide*, *Oracle8i Recovery Manager User's Guide and Reference*, and *Oracle8i Backup and Recovery Guide* for examples of performing media recovery.
- "[CREATE DATABASE](#)" on page 7-267 for information on creating a database.

**Prerequisites**

You must have ALTER DATABASE system privilege.

To specify the RECOVER clause, you must also have the OSDBA role enabled.

**Keywords and Parameters**


---

*database* identifies the database to be altered. The database name can contain only ASCII characters. If you omit database, Oracle alters the database identified by the value of the initialization parameter DB\_NAME. You can alter only the database whose control files are specified by the initialization parameter CONTROL\_FILES. The database identifier is not related to the Net8 database specification.

You can use the following clauses only when the database is not mounted by your instance:

MOUNT	mounts the database.
	STANDBY DATABASE      mounts the standby database. <b>See Also:</b> <i>Oracle8i Standby Database Concepts and Administration</i> .
	CLONE DATABASE      mounts the clone database. <b>See Also:</b> <i>Oracle8i Backup and Recovery Guide</i> .
CONVERT	completes the conversion of the Oracle7 data dictionary. After you use this clause, the Oracle7 data dictionary no longer exists in the Oracle database. <b>Use this clause only when you are migrating to Oracle8i.</b> <b>See Also:</b> <i>Oracle8i Migration</i> .
ACTIVATE STANDBY DATABASE	changes the state of a standby database to an active database. <b>See Also:</b> <i>Oracle8i Standby Database Concepts and Administration</i> .
OPEN	opens the database, making it available for normal use. You must mount the database before you can open it. You must activate a standby database before you can open it.

---

READ ONLY	restricts users to read-only transactions, preventing them from generating redo logs. You can use this clause to make a standby database available for queries even while archive logs are being copied from the primary database site.
	<p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot open a database READ ONLY if it is currently opened READ WRITE by another instance.</li> <li>■ You cannot open a database READ ONLY if it requires recovery.</li> <li>■ You cannot take tablespaces offline while the database is open READ ONLY. However, you can take datafiles offline and online, and you can recover offline datafiles and tablespaces while the database is open READ ONLY.</li> </ul>
READ WRITE	opens the database in read-write mode, allowing users to generate redo logs. This is the default.
RESETLOGS	resets the current log sequence number to 1 and discards any redo information that was not applied during recovery, ensuring that it will never be applied. This effectively discards all changes that are in the redo log, but not in the database.
	You must specify RESETLOGS to open the database after performing media recovery with an incomplete recovery using the RECOVER clause or with a backup control file. After opening the database with this clause, you should perform a complete database backup.
NORESETLOGS	leaves the log sequence number and redo log files in their current state.

**Restriction:** You can specify RESETLOGS and NORESETLOGS only after performing incomplete media recovery or complete media recovery with a backup control file. In any other case, Oracle uses the NORESETLOGS automatically.

You can use any of the following clauses when your instance has the database mounted, open or closed, and the files involved are not in use:

*general\_recovery\_clause* lets you design media recovery for the database or standby database, or for specified tablespaces or files.

**See Also:** *Oracle8i Backup and Recovery Guide* for more information on media recovery.

**Note:** If you do not have special media requirements, Oracle Corporation recommends that you use the SQL\*Plus RECOVER statement.

**See Also:** *SQL\*Plus User's Guide and Reference*.

---



**Restrictions:**

- You can recover the entire database only when the database is closed.
- Your instance must have the database mounted in exclusive mode.
- You can recover tablespaces or datafiles when the database is open or closed, provided that the tablespaces or datafiles to be recovered are offline.
- You cannot perform media recovery if you are connected to Oracle through the multi-threaded server architecture.

AUTOMATIC

automatically generates the name of the next archived redo log file needed to continue the recovery operation. If the `LOG_ARCHIVE_DEST_n` parameters are defined, Oracle scans those that are valid and enabled for the first local destination. It uses that destination in conjunction with `LOG_ARCHIVE_FORMAT` to generate the target redo log filename. If the `LOG_ARCHIVE_DEST_n` parameters are not defined, Oracle uses the value of the `LOG_ARCHIVE_DEST` parameter instead.

If the resulting file is found, Oracle applies the redo contained in that file. If the file is not found, Oracle prompts you for a filename, displaying the generated filename as a suggestion.

If you specify neither `AUTOMATIC` nor `LOGFILE`, Oracle prompts you for a filename, displaying the generated filename as a suggestion. You can then accept the generated filename or replace it with a fully qualified filename. If you know the archived filename differs from what Oracle would generate, you can save time by using the `LOGFILE` clause.

FROM  
'location'

specifies the location from which the archived redo log file group is read. The value of *location* must be a fully specified file location following the conventions of your operating system. If you omit this parameter, Oracle assumes the archived redo log file group is in the location specified by the initialization parameter `LOG_ARCHIVE_DEST` or `LOG_ARCHIVE_DEST_1`.

STANDBY  
DATABASE

recovers the standby database using the control file and archived redo log files copied from the primary database. The standby database must be mounted but not open.

DATABASE

recovers the entire database. This is the default. You can use this clause only when the database is closed.

---

**Note:** This clause recovers only online datafiles.

---

UNTIL

specifies the duration of the recovery operation.

- `CANCEL` performs cancel-based recovery. This clause recovers the database until you issue the `ALTER DATABASE RECOVER` statement with the `RECOVER CANCEL` clause.
- `TIME` performs time-based recovery. This parameter recovers the database to the time specified by the date. The date must be a character literal in the format 'YYYY-MM-DD:HH24:MI:SS'.

---

	<ul style="list-style-type: none"> <li>▪ <code>CHANGE</code> performs change-based recovery. This parameter recovers the database to a transaction-consistent state immediately before the system change number (SCN) specified by <i>integer</i>.</li> </ul>
	<p><code>USING BACKUP CONTROLFILE</code> specifies that a backup control file is being used instead of the current control file.</p>
<code>TABLESPACE</code>	recovers only the specified tablespaces. You can use this clause if the database is open or closed, provided the tablespaces to be recovered are offline.
<code>DATAFILE</code>	recovers the specified datafiles. You can use this clause when the database is open or closed, provided the datafiles to be recovered are offline.
<code>STANDBY TABLESPACE</code>   <code>DATAFILE</code>	reconstructs a lost or damaged datafile or tablespace in the standby database using archived redo log files copied from the primary database and a control file.
	<p><code>UNTIL</code> [ <code>CONSISTENT WITH</code> ] <code>CONTROLFILE</code> specifies that the recovery of an old standby datafile or tablespace uses the current standby database control file. However, any redo in advance of the standby controlfile will not be applied. The keywords <code>CONSISTENT WITH</code> are optional and are provided for semantic clarity.</p>
<code>LOGFILE</code>	continues media recovery by applying the specified redo log file.
<code>CONTINUE</code>	continues multi-instance recovery after it has been interrupted to disable a thread.
<code>CONTINUE DEFAULT</code>	continues recovery using the redo log file that Oracle would automatically generate if no other logfile were specified. This clause is equivalent to specifying <code>AUTOMATIC</code> , except that Oracle does not prompt for a filename.
<code>CANCEL</code>	terminates cancel-based recovery.
<i>managed_recovery_clause</i>	<p>specifies automated standby recovery mode. This mode assumes that the automated standby database is an active component of an overall standby database architecture. A primary database actively archives its redo log files to the standby site. As these archived redo logs arrive at the standby site, they become available for use by a managed standby recovery operation. Automated standby recovery is restricted to media recovery.</p> <p><b>See Also:</b> <i>Oracle8i Backup and Recovery Guide</i> for more information on the parameters of this clause.</p> <p><b>Restrictions:</b> The same restrictions apply as are listed under <a href="#">general_recovery_clause</a>.</p>
<code>TIMEOUT</code> <i>integer</i>	<p>specifies in minutes the wait period of the managed recovery operation. The recovery process waits for <i>integer</i> minutes for a requested archived log redo to be available for writing to the automated standby database. If the redo log file does not become available within that time, the recovery process terminates with an error message. You can then issue the statement again to return to automated standby recovery mode.</p> <p>If you do not specify this clause, the database remains in automated standby recovery mode until you reissue the statement with the <code>RECOVER CANCEL</code> clause or until instance shutdown or failure.</p>

---

CANCEL	terminates the managed recovery operation after applying all the redo in the current archived redo file.
CANCEL IMMEDIATE	terminates the managed recovery operation after applying all the redo in the current archived redo file or after the next redo log file read, whichever comes first. <b>Restriction:</b> This clause cannot be issued from the same session that issued the RECOVER MANAGED STANDBY DATABASE statement.
<i>parallel_clause</i>	specifies whether the recovery of media will be parallelized. For additional information, see the <a href="#">Notes</a> to the <i>parallel_clause</i> of "CREATE TABLE" on page 4-381.
NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the PARALLEL_THREADS_PER_CPU initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution processes. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .
RENAME GLOBAL_ NAME	changes the global name of the database. The <i>database</i> is the new database name and can be as long as eight bytes. The optional <i>domain</i> specifies where the database is effectively located in the network hierarchy. <b>See Also:</b> <i>Oracle8i Distributed Database Systems</i> for more information on global names. <b>Note:</b> Renaming your database does not change global references to your database from existing database links, synonyms, and stored procedures and functions on remote databases. Changing such references is the responsibility of the administrator of the remote databases.
RENAME FILE	renames datafiles, tempfiles, or redo log file members. This clause renames only files in the control file. It does not actually rename them on your operating system. You must specify each filename using the conventions for filenames on your operating system before specifying this clause.
RESET COMPATIBILITY	marks the database to be reset to an earlier version of Oracle when the database is next restarted. <b>Note:</b> RESET COMPATIBILITY works only if you have successfully disabled Oracle features that affect backward compatibility. <b>See Also:</b> <i>Oracle8i Migration</i> for more information on downgrading to an earlier version of Oracle.

---

You can use the following clauses only when your instance has the database open:

ENABLE THREAD	in a parallel server, enables the specified thread of redo log file groups. The thread must have at least two redo log file groups before you can enable it.
---------------	--

---

	PUBLIC	makes the enabled thread available to any instance that does not explicitly request a specific thread with the initialization parameter <code>THREAD</code> . If you omit <code>PUBLIC</code> , the thread is available only to the instance that explicitly requests it with the initialization parameter <code>THREAD</code> .
DISABLE	THREAD	disables the specified thread, making it unavailable to all instances. You cannot disable a thread if an instance using it has the database mounted.

**See Also:** *Oracle8i Designing and Tuning for Performance* for more information on enabling and disabling threads.

CHARACTER SET	CHARACTER SET	changes the character set the database uses to store data. <code>NATIONAL CHARACTER SET</code> changes the national character set used to store data in columns specifically defined as <code>NCHAR</code> , <code>NCLOB</code> , or <code>NVARCHAR2</code> . Specify <i>character_set</i> without quotation marks.
NATIONAL CHARACTER SET		

---

**WARNING:** You cannot roll back an `ALTER DATABASE CHARACTER SET` or `ALTER DATABASE NATIONAL CHARACTER SET` statement. Therefore, you should perform a full backup before issuing either of these statements.

---

**Restrictions:**

- You must have `SYSDBA` system privilege, and you must start up the database in restricted mode (for example, with the `SQL*Plus STARTUP RESTRICT` command).
- The current character set must be a strict subset of the character set to which you change. That is, each character represented by a codepoint value in the source character set must be represented by the same codepoint value in the target character set. For a list of valid character sets, see *Oracle8i National Language Support Guide*.

*datafile/tempfile\_*  
*clauses* let you modify datafiles and tempfiles.

You can use any of the following clauses when your instance has the database mounted, open or closed, and the files involved are not in use:

CREATE DATAFILE	creates a new empty datafile in place of an old one. You can use this clause to re-create a datafile that was lost with no backup. The ' <i>filename</i> ' must identify a file that is or was once part of the database. The <i>filespec</i> specifies the name and size of the new datafile. If you omit the <code>AS</code> clause, Oracle creates the new file with the name and size as the file specified by ' <i>filename</i> '.
--------------------	---

During recovery, all archived redo logs written to since the original datafile was created must be applied to the new, empty version of the lost datafile.

Oracle creates the new file in the same state as the old file when it was created. You must perform media recovery on the new file to return it to the state of the old file at the time it was lost.

**Restriction:** You cannot create a new file based on the first datafile of the `SYSTEM` tablespace.

DATAFILE	affects your database files as follows:
----------	---

---

ONLINE	brings the datafile online.
OFFLINE	takes the datafile offline. If the database is open, you must perform media recovery on the datafile before bringing it back online, because a checkpoint is not performed on the datafile before it is taken offline.  DROP takes a datafile offline when the database is in NOARCHIVELOG mode.
RESIZE	attempts to increase or decrease the size of the datafile to the specified absolute size in bytes. Use K or M to specify this size in kilobytes or megabytes. There is no default, so you must specify a size.  If sufficient disk space is not available for the increased size, or if the file contains data beyond the specified decreased size, Oracle returns an error.
<i>autoextend_</i> <i>clause</i>	enables or disables the automatic extension of a datafile. If you do not specify this clause, datafiles are not automatically extended.  OFF disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in further ALTER DATABASE AUTOEXTEND statements.  ON enables autoextend.  NEXT specifies in bytes the size of the next increment of disk space to be automatically allocated to the datafile when more extents are required. Use K or M to specify this size in kilobytes or megabytes. The default is one data block.  MAXSIZE specifies the maximum disk space allowed for automatic extension of the datafile.  UNLIMITED sets no limit on allocating disk space to the datafile.
END BACKUP	avoids media recovery on database startup after an online tablespace backup was interrupted by a system failure or instance failure or SHUTDOWN ABORT.
<hr/>	
<b>WARNING: Do not use ALTER TABLESPACE ... END BACKUP if you have restored any of the files affected from a backup. Media recovery is fully described in <i>Oracle8i Backup and Recovery Guide</i>.</b>	
TEMPFILE	Lets you resize your temporary datafile or specify the <i>autoextend_clause</i> , with the same effect as with a permanent datafile.  <b>Restriction:</b> You cannot specify TEMPFILE unless the database is open.
DROP	drops <i>tempfile</i> from the database. The tablespace remains.
<i>logfile_clauses</i>	lets you add, drop, or modify log files.

**ARCHIVELOG** specifies that the contents of a redo log file group must be archived before the group can be reused. This mode prepares for the possibility of media recovery. Use this clause only after shutting down your instance normally or immediately with no errors and then restarting it, mounting the database in parallel server disabled mode.

**NOARCHIVELOG** specifies that the contents of a redo log file group need not be archived so that the group can be reused. This mode does not prepare for recovery after media failure.

Use the **ARCHIVELOG** clause and **NOARCHIVELOG** clause only if your instance has the database mounted in parallel server disabled mode, but not open.

**ADD LOGFILE** adds one or more redo log file groups to the specified thread, making them available to the instance assigned the thread.

**THREAD *integer*** is applicable only if you are using Oracle with the Parallel Server option in parallel mode. *integer* is the thread number. The number of threads you can create is limited by the value of the **MAXINSTANCES** parameter specified in the **CREATE DATABASE** statement.

If you omit **THREAD**, the redo log file group is added to the thread assigned to your instance.

**GROUP *integer*** uniquely identifies the redo log file group among all groups in all threads and can range from 1 to the **MAXLOGFILES** value. You cannot add multiple redo log file groups having the same **GROUP** value. If you omit this parameter, Oracle generates its value automatically. You can examine the **GROUP** value for a redo log file group through the dynamic performance view **V\$LOG**.

***filespec*** Each *filespec* specifies a redo log file group containing one or more members, or copies. See the syntax description of *filespec* in "[filespec](#)" on page 7-516.

**ADD LOGFILE MEMBER** adds new members to existing redo log file groups. Each new member is specified by '*filename*'. If the file already exists, it must be the same size as the other group members, and you must specify **REUSE**. If the file does not exist, Oracle creates a file of the correct size. You cannot add a member to a group if all of the group's members have been lost through media failure.

You can specify an existing redo log file group in one of these ways:

**GROUP *integer*** Specify the value of the **GROUP** parameter that identifies the redo log file group.

***list of filenames*** List all members of the redo log file group. You must fully specify each filename according to the conventions of your operating system.

---

DROP LOGFILE	<p>drops all members of a redo log file group. Specify a redo log file group as indicated for the <code>ADD LOGFILE MEMBER</code> clause.</p> <ul style="list-style-type: none"> <li>■ To drop the current log file group, you must first issue an <code>ALTER SYSTEM SWITCH LOGFILE</code> statement. See "<a href="#">ALTER SYSTEM</a>" on page 7-102.</li> <li>■ You cannot drop a redo log file group if it needs archiving.</li> <li>■ You cannot drop a redo log file group if doing so would cause the redo thread to contain less than two redo log file groups.</li> </ul>
DROP LOGFILE MEMBER	<p>drops one or more redo log file members. Each 'filename' must fully specify a member using the conventions for filenames on your operating system.</p> <ul style="list-style-type: none"> <li>■ To drop a log file in the current log, you must first issue an <code>ALTER SYSTEM SWITCH LOGFILE</code> statement. See "<a href="#">ALTER SYSTEM</a>" on page 7-102.</li> <li>■ You cannot use this clause to drop all members of a redo log file group that contains valid data. To perform this operation, use the <code>DROP LOGFILE</code> clause.</li> </ul>
CLEAR LOGFILE	<p>reinitializes an online redo log, optionally without archiving the redo log. <code>CLEAR LOGFILE</code> is similar to adding and dropping a redo log, except that the statement may be issued even if there are only two logs for the thread and also may be issued for the current redo log of a closed thread.</p> <p><code>UNARCHIVED</code>      You must specify <code>UNARCHIVED</code> if you want to reuse a redo log that was not archived.</p> <p style="text-align: center;"><b>WARNING: Specifying <code>UNARCHIVED</code> makes backups unusable if the redo log is needed for recovery.</b></p> <p>Do not use <code>CLEAR LOGFILE</code> to clear a log needed for media recovery. If it is necessary to clear a log containing redo after the database checkpoint, you must first perform incomplete media recovery. The current redo log of an open thread can be cleared. The current log of a closed thread can be cleared by switching logs in the closed thread.</p> <p>If the <code>CLEAR LOGFILE</code> statement is interrupted by a system or instance failure, then the database may hang. If this occurs, reissue the statement after the database is restarted. If the failure occurred because of I/O errors accessing one member of a log group, then that member can be dropped and other members added.</p> <p><code>UNRECOVER- ABLE DATAFILE</code>      You must specify <code>UNRECOVERABLE DATAFILE</code> if you have taken the datafile offline with the database in <code>ARCHIVELOG</code> mode (that is, you specified <code>ALTER DATABASE . . . DATAFILE OFFLINE</code> without the <code>DROP</code> keyword), and if the unarchived log to be cleared is needed to recover the datafile before bringing it back online. In this case, you must drop the datafile and the entire tablespace once the <code>CLEAR LOGFILE</code> statement completes.</p>

*controlfile\_clauses*

CREATE STANDBY CONTROLFILE	creates a control file to be used to maintain a standby database. If the file already exists, you must specify <code>REUSE</code> .  <b>See Also:</b> <i>Oracle8i Standby Database Concepts and Administration</i> .
BACKUP CONTROLFILE	backs up the current control file.  TO ' <i>filename</i> ' specifies the file to which the control file is backed up. You must fully specify the <i>filename</i> using the conventions for your operating system. If the specified file already exists, you must specify <code>REUSE</code> .  TO TRACE writes SQL statements to the database's trace file rather than making a physical backup of the control file. The SQL statements can start up the database, re-create the control file, and recover and open the database appropriately, based on the created control file.  You can copy the statements from the trace file into a script file, edit the statements as necessary, and use the database if all copies of the control file are lost (or to change the size of the control file). <ul style="list-style-type: none"><li>■ <code>RESETLOGS</code> specifies that the SQL statement written to the trace file for starting the database is <code>ALTER DATABASE OPEN RESETLOGS</code>.</li><li>■ <code>NORESETLOGS</code> specifies that the SQL statement written to the trace file for starting the database is <code>ALTER DATABASE OPEN NORESETLOGS</code>.</li></ul>

---

## Examples

**READ ONLY / READ WRITE Example** The first statement below opens the database in read-only mode. The second statement returns the database to read-write mode and clears the online redo logs:

```
ALTER DATABASE OPEN READ ONLY;
```

```
ALTER DATABASE OPEN READ WRITE RESETLOGS;
```

**PARALLEL Example** The following statement performs tablespace recovery using parallel recovery processes:

```
ALTER DATABASE  
  RECOVER TABLESPACE binky  
  PARALLEL;
```

**Redo Log File Group Example** The following statement adds a redo log file group with two members and identifies it with a `GROUP` parameter value of 3:



```
ALTER DATABASE stocks
  ADD LOGFILE GROUP 3
    ('diska:log3.log' ,
     'diskb:log3.log') SIZE 50K;
```

**Redo Log File Group Member Example** The following statement adds a member to the redo log file group added in the previous example:

```
ALTER DATABASE stocks
  ADD LOGFILE MEMBER 'diskc:log3.log'
  TO GROUP 3;
```

**Dropping a Log File Member** The following statement drops the redo log file member added in the previous example:

```
ALTER DATABASE stocks
  DROP LOGFILE MEMBER 'diskc:log3.log';
```

**Renaming a Log File Member** The following statement renames a redo log file member:

```
ALTER DATABASE stocks
  RENAME FILE 'diskb:log3.log' TO 'diskd:log3.log';
```

The above statement only changes the member of the redo log group from one file to another. The statement does not actually change the name of the file 'DISKB:LOG3.LOG' to 'DISKD:LOG3.LOG'. You must perform this operation through your operating system.

**Dropping All Log File Group Members** The following statement drops all members of the redo log file group 3:

```
ALTER DATABASE stocks DROP LOGFILE GROUP 3;
```

**Adding a Redo Log File Group** The following statement adds a redo log file group containing three members to thread 5 (in an Oracle Parallel Server environment) and assigns it a GROUP parameter value of 4:

```
ALTER DATABASE stocks
  ADD LOGFILE THREAD 5 GROUP 4
    ('diska:log4.log' ,
     'diskb:log4.log' ,
     'diskc:log4.log' );
```

**Disabling a Parallel Server Thread** The following statement disables thread 5 in a parallel server:

```
ALTER DATABASE stocks
  DISABLE THREAD 5;
```

**Enabling a Parallel Server Thread** The following statement enables thread 5 in a parallel server, making it available to any Oracle instance that does not explicitly request a specific thread:

```
ALTER DATABASE stocks
  ENABLE PUBLIC THREAD 5;
```

**Creating a New Datafile** The following statement creates a new datafile 'DISK2:DB1.DAT' based on the file 'DISK1:DB1.DAT':

```
ALTER DATABASE
  CREATE DATAFILE 'disk1:db1.dat' AS 'disk2:db1.dat';
```

**Changing the Global Database Name** The following statement changes the global name of the database and includes both the database name and domain:

```
ALTER DATABASE
  RENAME GLOBAL_NAME TO sales.australia.acme.com;
```

**Character Set Example** The following statements change the database character set and national character set to the WE8ISO8859P1 character set:

```
ALTER DATABASE db1 CHARACTER SET WE8ISO8859P1;
ALTER DATABASE db1 NATIONAL CHARACTER SET WE8ISO8859P1;
```

The database name is optional, and the character set name is specified without quotation marks.

**Resizing a Datafile** The following statement attempts to change the size of datafile 'DISK1:DB1.DAT':

```
ALTER DATABASE
  DATAFILE 'disk1:db1.dat' RESIZE 10 M;
```

**Clearing a Log File** The following statement clears a log file:

```
ALTER DATABASE
  CLEAR LOGFILE 'disk3:log.dbf';
```

**Database Recovery Examples** The following statement performs complete recovery of the entire database, letting Oracle generate the name of the next archived redo log file needed:

```
ALTER DATABASE
  RECOVER AUTOMATIC DATABASE;
```

The following statement explicitly names a redo log file for Oracle to apply:

```
ALTER DATABASE
  RECOVER LOGFILE 'diska:arch0006.arc';
```

The following statement performs time-based recovery of the database:

```
ALTER DATABASE
  RECOVER AUTOMATIC UNTIL TIME '1998-10-27:14:00:00';
```

Oracle recovers the database until 2:00 pm on October 27, 1998.

The following statement recovers the tablespace USER5:

```
ALTER DATABASE
  RECOVER TABLESPACE user5;
```

The following statement recovers the standby datafile /FINANCE/STBS\_21.f, using the corresponding datafile in the original standby database, plus all relevant archived logs and the current standby database control file:

```
ALTER DATABASE
  RECOVER STANDBY DATAFILE '/finance/stbs_21.f'
  UNTIL CONTROLFILE;
```

**Managed Standby Database Examples** The following statement recovers the standby database in automated standby recovery mode:

```
ALTER DATABASE
  RECOVER MANAGED STANDBY DATABASE;
```

The following statement puts the database in automated standby recovery mode. The managed recovery process will wait up to 60 minutes for the next archive log:

```
ALTER DATABASE
  RECOVER MANAGED STANDBY DATABASE TIMEOUT 60;
```

If each subsequent log arrives within 60 minutes of the last log, recovery continues indefinitely or until manually terminated.

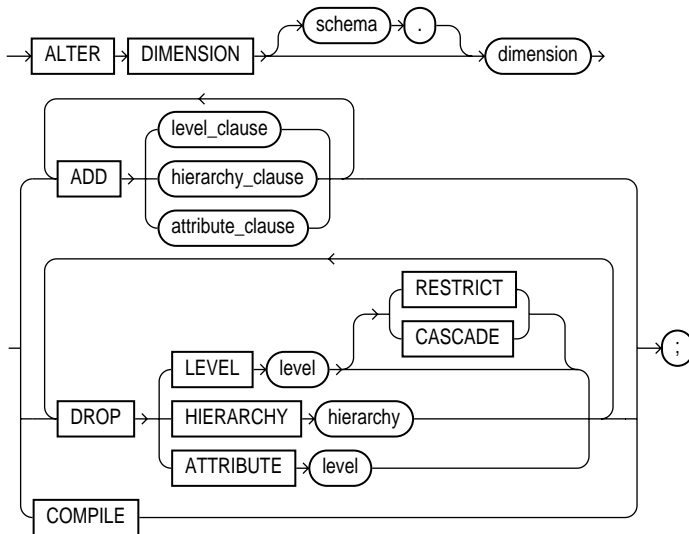
The following statement terminates the managed recovery operation:

```
ALTER DATABASE  
    RECOVER MANAGED STANDBY DATABASE CANCEL IMMEDIATE;
```


The managed recovery operation terminates before the next group of redo is read from the current redo log file. Media recovery ends in the "middle" of applying redo from the current redo log file.

# ALTER DIMENSION


## Syntax




### level\_clause::=




### hierarchy\_clause::=



**join\_clause::=**



**attribute\_clause::=**



## Purpose

To change the hierarchical relationships or dimension attributes of a dimension.

**See Also:** ["CREATE DIMENSION"](#) on page 7-277 for more information on dimensions,

## Prerequisites

The dimension must be in your schema or you must have the ALTER ANY DIMENSION system privilege to use this statement.

A dimension is always altered under the rights of the owner.

## Keywords and Parameters

The following keywords and parameters have meaning unique to ALTER DIMENSION. The remaining keywords and parameters have the same functionality that they have in the CREATE DIMENSION statement. See ["CREATE DIMENSION"](#) on page 7-277.

---

<i>schema</i>	is the schema of the dimension you want to modify. If you do not specify <i>schema</i> , Oracle assumes the dimension is in your own schema.
<i>dimension</i>	is the name of the dimension. This dimension must already exist.

ADD	lets you add a level, hierarchy, or attribute to the dimension. Adding one of these elements does not invalidate any existing materialized view.  Oracle processes ADD LEVEL clauses prior to any other ADD clauses.
DROP	lets you drop a level, hierarchy, or attribute from the dimension. Any level, hierarchy, or attribute you specify must already exist.  <b>Restriction:</b> If any attributes or hierarchies reference a level, you cannot drop the level until you either drop all the referencing attributes and hierarchies or specify CASCADE.  CASCADE causes Oracle to drop any attributes or hierarchies that reference the level, along with the level itself.  RESTRICT prevents Oracle from dropping a level that is referenced by any attributes or hierarchies. This is the default.
COMPILE	explicitly recompiles an invalidated dimension. Oracle automatically compiles a dimension when you issue an ADD clause or DROP clause. However, if you alter an object referenced by the dimension (for example, if you drop and then re-create a table referenced in the dimension), the dimension will be invalidated, and you must recompile it explicitly.

---

## Examples


This example modifies the TIME dimension:

```
ALTER DIMENSION time
  DROP HIERARCHY week_month;
ALTER DIMENSION time
  DROP ATTRIBUTE cur_date;
ALTER DIMENSION time
  ADD LEVEL day IS time_tab.t_day
  ADD ATTRIBUTE day DETERMINES t_holiday;
```

---

## ALTER FUNCTION

### Syntax



### Purpose

To recompile an invalid standalone stored function. Explicit recompilation eliminates the need for implicit run-time recompilation and prevents associated run-time compilation errors and performance overhead.

The `ALTER FUNCTION` statement is similar to "[ALTER PROCEDURE](#)" on page 2-67. For information on how Oracle recompiles functions and procedures, see *Oracle8i Concepts*.

---

**Note:** This statement does not change the declaration or definition of an existing function. To redeclare or redefine a function, use the `CREATE FUNCTION` statement with the `OR REPLACE` clause; see "[CREATE FUNCTION](#)" on page 7-284.

---

### Prerequisites

The function must be in your own schema or you must have `ALTER ANY PROCEDURE` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the function. If you omit <i>schema</i> , Oracle assumes the function is in your own schema.
<i>function</i>	is the name of the function to be recompiled.
<code>COMPILE</code>	causes Oracle to recompile the function. The <code>COMPILE</code> keyword is required. If Oracle does not compile the function successfully, you can see the associated compiler error messages with the SQL*Plus command <code>SHOW ERRORS</code> .
<code>DEBUG</code>	instructs the PL/SQL compiler to generate and store the code for use by the PL/SQL debugger.

---



## Example

To explicitly recompile the function `GET_BAL` owned by the user `MERRIWEATHER`, issue the following statement:

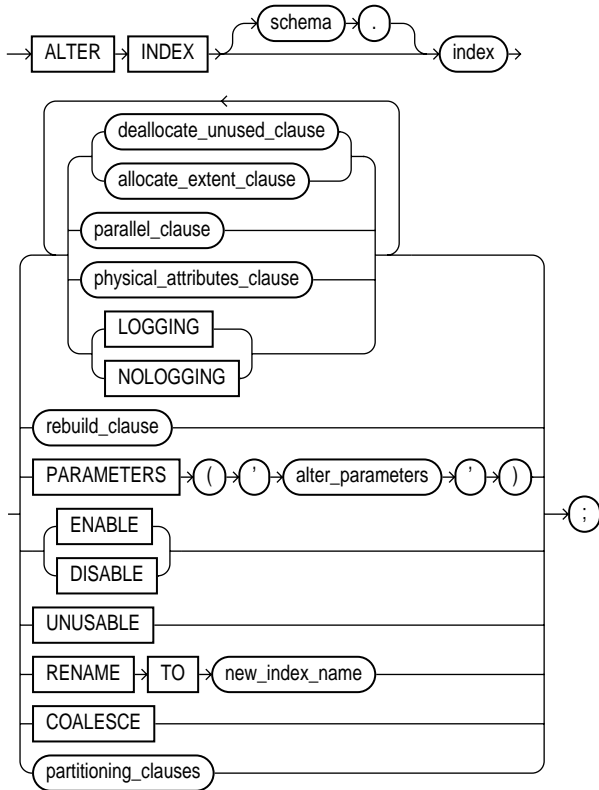
```
ALTER FUNCTION merriweather.get_bal  
    COMPILE;
```

If Oracle encounters no compilation errors while recompiling `GET_BAL`, `GET_BAL` becomes valid. Oracle can subsequently execute it without recompiling it at run time. If recompiling `GET_BAL` results in compilation errors, Oracle returns an error, and `GET_BAL` remains invalid.


Oracle also invalidates all objects that depend upon `GET_BAL`. If you subsequently reference one of these objects without explicitly recompiling it first, Oracle recompiles it implicitly at run time.

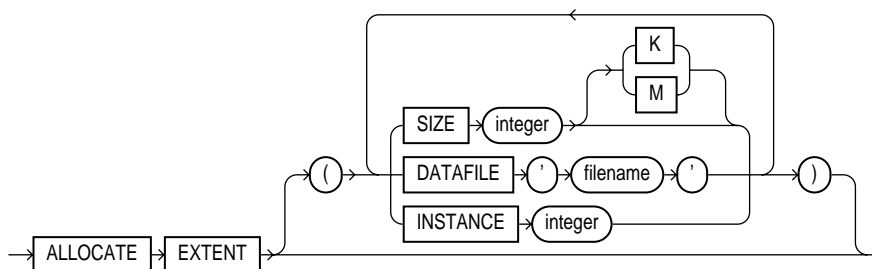


# ALTER INDEX

## Syntax



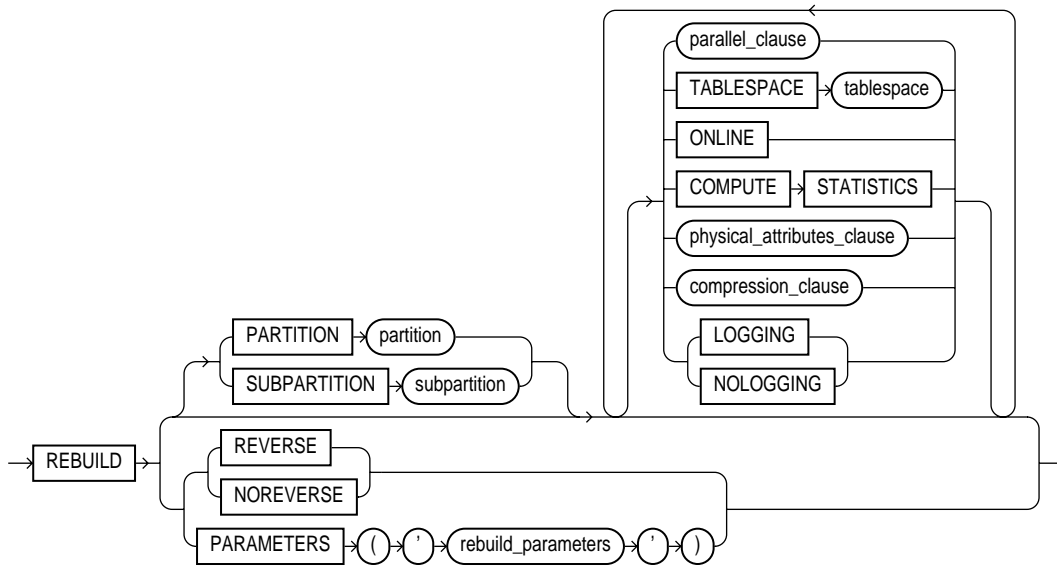
### deallocate\_unused\_clause::=




**allocate\_extent\_clause::=****parallel\_clause::=****physical\_attributes\_clause::=**

**storage\_clause:** See "storage\_clause" on page 7-605.

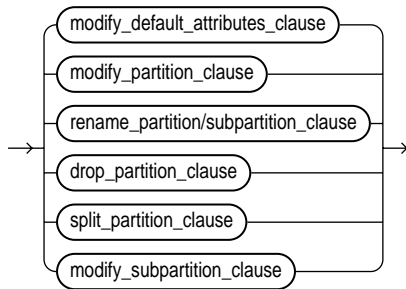
**rebuild\_clause::=**

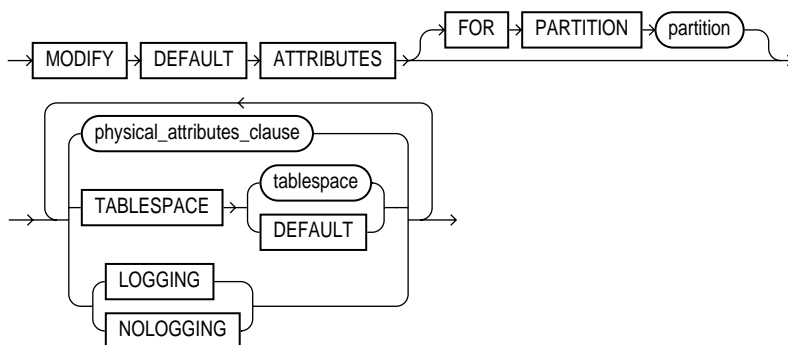
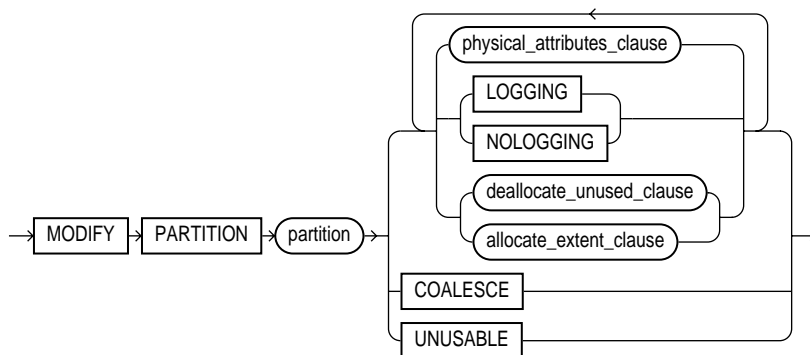






**compression\_clause::=**



**partitioning\_clauses::=**



**modify\_default\_attributes\_clause::=****modify\_partition\_clause::=****rename\_partition/ subpartition\_clause::=****drop\_partition\_clause::=**

**split\_partition\_clause::=****partition\_description::=****modify\_subpartition\_clause::=****Purpose**

To change or rebuild an existing index.

For information on creating an index, see "[CREATE INDEX](#)" on page 7-291.

**Prerequisites**

The index must be in your own schema or you must have `ALTER ANY INDEX` system privilege.

Schema object privileges are granted on the parent index, not on individual index partitions or subpartitions.

You must have tablespace quota to modify, rebuild, or split an index partition or to modify or rebuild an index subpartition.

## Keywords and Parameters

<i>schema</i>	is the schema containing the index. If you omit <i>schema</i> , Oracle assumes the index is in your own schema.
<i>index</i>	is the name of the index to be altered.  <b>Restrictions:</b> <ul style="list-style-type: none"> <li>■ If <i>index</i> is a domain index, you can specify only the <code>PARAMETERS</code> clause, the <code>RENAME</code> clause, or the <i>rebuild_clause</i> (with or without the <code>PARAMETERS</code> clause). No other clauses are valid.</li> <li>■ You cannot alter or rename a domain index that is marked <code>LOADING</code> or <code>FAILED</code>. If an index is marked <code>FAILED</code>, the only clause you can specify is <code>REBUILD</code>. For information on the <code>LOADING</code> and <code>FAILED</code> states of domain indexes, see <i>Oracle8i Data Cartridge Developer's Guide</i>.</li> </ul>
<i>deallocate_unused_clause</i>	explicitly deallocates unused space at the end of the index and makes the freed space available for other segments in the tablespace. Only unused space above the high water mark can be freed.  If <i>index</i> is range-partitioned or hash-partitioned, Oracle deallocates unused space from each index partition. If <i>index</i> is a local index on a composite-partitioned table, Oracle deallocates unused space from each index subpartition.  <b>Restrictions:</b> <ul style="list-style-type: none"> <li>■ You cannot specify this clause for an index on a temporary table.</li> <li>■ You cannot specify this clause and also specify the <i>rebuild_clause</i>.</li> </ul> <p><b>See Also:</b> "<a href="#">ALTER TABLE</a>" on page 7-123 for more information on this clause.</p>
<code>KEEP</code>	specifies the number of bytes above the high water mark that the index will have after deallocation. If the number of remaining extents are less than <code>MINEXTENTS</code> , then <code>MINEXTENTS</code> is set to the current number of extents. If the initial extent becomes smaller than <code>INITIAL</code> , then <code>INITIAL</code> is set to the value of the current initial extent. If you omit <code>KEEP</code> , all unused space is freed.  <b>See Also:</b> " <a href="#">ALTER TABLE</a> " on page 7-123 for a complete description of this clause.
<i>allocate_extent_clause</i>	explicitly allocates a new extent for the index. For a local index on a hash-partitioned table, Oracle allocates a new extent for each partition of the index.  <b>Restriction:</b> You cannot specify this clause for an index on a temporary table or for a range-partitioned or composite-partitioned index.
<code>SIZE</code>	specifies the size of the extent in bytes. Use <code>K</code> or <code>M</code> to specify the extent size in kilobytes or megabytes. If you omit <code>SIZE</code> , Oracle determines the size based on the values of the index's storage parameters.

---

DATAFILE	specifies one of the datafiles in the index's tablespace to contain the new extent. If you omit DATAFILE, Oracle chooses the datafile.
INSTANCE	makes the new extent available to the specified instance. An instance is identified by the value of its initialization parameter INSTANCE_NUMBER. If you omit this parameter, the extent is available to all instances. Use this parameter only if you are using Oracle with the Parallel Server option in parallel mode.

Explicitly allocating an extent with this clause does not change the values of the NEXT and PCTINCREASE storage parameters, so does not affect the size of the next extent to be allocated.

*parallel\_clause*

changes the default degree of parallelism for queries and DML on the index. For additional information, see the [Notes](#) to the [parallel\\_clause](#) of "CREATE TABLE" on page 4-381.

NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances multiplied by the value of the PARALLEL_THREADS_PER_CPU initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution processes. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .

**Restriction:** You cannot specify this clause for an index on a temporary table.

*physical\_attributes\_clause*

lets you change the values of parameters for a nonpartitioned index, all partitions and subpartitions of a partitioned index, a specified partition, or all subpartitions of a specified partition. See these parameters in "CREATE TABLE" on page 4-381.

**Restrictions:**

- You cannot specify this clause for an index on a temporary table.
- You cannot specify the PCTUSED parameter when altering an index.
- You cannot change the value of the PCTFREE parameter for the index as a whole (ALTER INDEX) or for a partition (ALTER INDEX ... MODIFY PARTITION). You can specify PCTFREE in all other forms of the ALTER INDEX statement.

*storage\_clause* changes the storage parameters for a nonpartitioned index, index partition, or all partitions of a partitioned index, or default values of these parameters for a partitioned index. See the ["storage\\_clause"](#) on page 7-605.

LOGGING |  
NOLOGGING

LOGGING | NOLOGGING specifies that subsequent Direct Loader (SQL\*Loader) and direct-load INSERT operations against a nonpartitioned index, a range or hash index partition, or all partitions or subpartitions of a composite-partitioned index will be logged (LOGGING) or not logged (NOLOGGING) in the redo log file.



In `NOLOGGING` mode, data is modified with minimal logging (to mark new extents invalid and to record dictionary changes). When applied during media recovery, the extent invalidation records mark a range of blocks as logically corrupt, because the redo data is not logged. Therefore, if you cannot afford to lose this index, you must take a backup after the operation in `NOLOGGING` mode.

If the database is run in `ARCHIVELOG` mode, media recovery from a backup taken before an operation in `LOGGING` mode will re-create the index. However, media recovery from a backup taken before an operation in `NOLOGGING` mode will not re-create the index.

An index segment can have logging attributes different from those of the base table and different from those of other index segments for the same base table.

**Restriction:** You cannot specify this clause for an index on a temporary table.

**See Also:** *Oracle8i Concepts* and the *Oracle8i Parallel Server Concepts* for more information about `LOGGING` and parallel DML.

`RECOVERABLE` |  
`UNRECOVERABLE`

These keywords are deprecated and have been replaced with `LOGGING` and `NOLOGGING`, respectively. Although `RECOVERABLE` and `UNRECOVERABLE` are supported for backward compatibility, Oracle Corporation strongly recommends that you use the `LOGGING` and `NOLOGGING` keywords.

`RECOVERABLE` is not a valid keyword for creating partitioned tables or LOB storage characteristics. `UNRECOVERABLE` is not a valid keyword for creating partitioned or index-organized tables. Also, it can be specified only with the `AS` subquery clause of `CREATE INDEX`.

*rebuild\_clause*

re-creates an existing index or one of its partitions or subpartitions. For a function-based index, this clause also enables the index. If the function on which the index is based does not exist, the rebuild statement will fail.

**Restrictions:**

- You cannot rebuild an index on a temporary table.
- You cannot rebuild an entire partitioned index. You must rebuild each partition or subpartition, as described below.
- You cannot also specify the *deallocate\_unused\_clause* in this statement.
- You cannot change the value of the `PCTFREE` parameter for the index as a whole (`ALTER INDEX`) or for a partition (`ALTER INDEX ... MODIFY PARTITION`). You can specify `PCTFREE` in all other forms of the `ALTER INDEX` statement.

<code>PARTITION <i>partition</i></code>	<p>rebuilds one partition of an index. You can also use this clause to move an index partition to another tablespace or to change a create-time physical attribute.</p> <p><b>Restriction:</b> You cannot specify this clause for a local index on a composite-partitioned table. Instead, use the <code>REBUILD SUBPARTITION</code> clause.</p> <p><b>See Also:</b> <i>Oracle8i Administrator's Guide</i> for more information about partition maintenance operations.</p>
<code>SUBPARTITION <i>subpartition</i></code>	<p>rebuilds one subpartition of an index. You can also use this clause to move an index subpartition to another tablespace. If you do not specify <code>TABLESPACE</code>, the subpartition is rebuilt in the same tablespace.</p> <p><b>Restrictions:</b> The only parameters you can specify for a subpartition are <code>TABLESPACE</code> and the <i>parallel_clause</i>.</p>
<code>REVERSE   NOEVERSE</code>	<p>specifies whether the bytes of the index block are stored in reverse order.</p> <ul style="list-style-type: none"><li>■ <code>REVERSE</code> stores the bytes of the index block in reverse order and excludes the rowid when the index is rebuilt.</li><li>■ <code>NOEVERSE</code> stores the bytes of the index block without reversing the order when the index is rebuilt. Rebuilding a <code>REVERSE</code> index without the <code>NOEVERSE</code> keyword produces a rebuilt, reverse-keyed index.</li></ul> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot reverse a bitmap index or an index-organized table.</li><li>■ You cannot specify <code>REVERSE</code> or <code>NOEVERSE</code> for a partition or subpartition.</li></ul>
<code>TABLESPACE</code>	<p>specifies the tablespace where the rebuilt index, index partition, or index subpartition will be stored. The default is the default tablespace where the index or partition resided before you rebuilt it.</p>

---

COMPRESS	<p>enables key compression, which eliminates repeated occurrence of key column values. Use <i>integer</i> to specify the prefix length (number of prefix columns to compress).</p> <ul style="list-style-type: none"> <li>■ For unique indexes, the range of valid prefix length values is from 1 to the number of key columns minus 1. The default prefix length is the number of key columns minus 1.</li> <li>■ For nonunique indexes, the range of valid prefix length values is from 1 to the number of key columns. The default prefix length is number of key columns.</li> </ul> <p>Oracle compresses only nonpartitioned indexes that are nonunique or unique indexes of at least two columns.</p> <p><b>Restriction:</b> You cannot specify COMPRESS for a bitmapped index.</p>
NOCOMPRESS	<p>disables key compression. This is the default.</p>
ONLINE	<p>specifies that DML operations on the table or partition are allowed during rebuilding of the index.</p> <p><b>Restriction:</b> Parallel DML is not supported during online index building. If you specify ONLINE and then issue parallel DML statements, Oracle returns an error.</p>
COMPUTE STATISTICS	<p>enables you to collect statistics at relatively little cost during the rebuilding of an index. These statistics are stored in the data dictionary for ongoing use by the optimizer in choosing a plan of execution for SQL statements.</p> <p>The types of statistics collected depend on the type of index you are rebuilding.</p> <hr/> <p><b>Note:</b> If you create an index using another index (instead of a table), the original index might not provide adequate statistical information. Therefore, Oracle generally uses the base table to compute the statistics, which will improve the statistics but may negatively affect performance.</p> <hr/> <p>Additional methods of collecting statistics are available in PL/SQL packages and procedures. See <i>Oracle8i Supplied PL/SQL Packages Reference</i>.</p>
LOGGING   NOLOGGING	<p>specifies whether the ALTER INDEX ... REBUILD operation will be logged.</p>

PARAMETERS	<p>applies only to <b>domain indexes</b>. This clause specifies the parameter string for altering the index (or, in the <i>rebuild_clause</i>, rebuilding the index). The maximum length of the parameter string is 1000 characters. This string is passed uninterpreted to the appropriate indextype routine.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"><li>■ <i>Oracle8i Data Cartridge Developer's Guide</i> for more information on these routines.</li><li>■ <a href="#">"CREATE INDEX"</a> on page 7-291 for more information on domain indexes.</li></ul> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot specify this clause for any indexes other than domain indexes.</li><li>■ The parameter string is passed to the appropriate routine only if <i>index</i> is not marked UNUSABLE.</li></ul>
ENABLE	<p>applies only to a <b>function-based index</b> that has been disabled because a user-defined function used by the index was dropped or replaced. This clause enables such an index if</p> <ul style="list-style-type: none"><li>■ the function is currently valid,</li><li>■ the signature of the current function matches the signature of the function when the index was created, and</li><li>■ the function is currently marked as DETERMINISTIC.</li></ul> <p><b>Restriction:</b> You cannot specify any other clauses of ALTER INDEX in the same statement with ENABLE.</p>
DISABLE	<p>applies only to a <b>function-based index</b>. This clause enables you to disable the use of a function-based index. You might want to do so, for example, while working on the body of the function. Afterward you can either rebuild the index or specify another ALTER INDEX statement with the ENABLE keyword.</p>
UNUSABLE	<p>marks the index or index partition(s) or index subpartition(s) UNUSABLE. An unusable index must be rebuilt, or dropped and re-created, before it can be used. While one partition is marked UNUSABLE, the other partitions of the index are still valid. You can execute statements that require the index if the statements do not access the unusable partition. You can also split or rename the unusable partition before rebuilding it.</p> <p><b>Restriction:</b> You cannot specify this clause for an index on a temporary table.</p>
RENAME TO	<p>renames <i>index</i> to <i>new_index_name</i>. The <i>new_index_name</i> is a single identifier and does not include the schema name.</p>
COALESCE	<p>instructs Oracle to merge the contents of index blocks where possible to free blocks for reuse.</p> <p><b>Restriction:</b> You cannot specify this clause for an index on a temporary table.</p> <p><b>See Also:</b> <i>Oracle8i Administrator's Guide</i> for more information on space management and coalescing indexes.</p>

*partitioning\_clauses:* The remainder of the clauses of the ALTER INDEX statement are valid only for partitioned indexes.

---

	<p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify any of these clauses for an index on a temporary table.</li> <li>■ You can combine several operations on the base index into one ALTER INDEX statement (except RENAME and REBUILD), but you cannot combine partition operations with other partition operations or with operations on the base index.</li> </ul>
<i>modify_default_attributes_clause</i>	<p>specifies new values for the default attributes of a partitioned index.</p> <p><b>Restriction:</b> The only attribute you can specify for an index on a hash-partitioned or composite-partitioned table is TABLESPACE.</p> <p>TABLESPACE      specifies the default tablespace for new partitions of an index or subpartitions of an index partition.</p> <p>LOGGING   NOLOGGING      specifies the default logging attribute of a partitioned index or an index partition.</p> <p>FOR PARTITION <i>partition</i>      specifies the default attributes for the subpartitions of a partition of a local index on a composite-partitioned table.</p>
<i>modify_partition_clause</i>	<p>modifies the real physical attributes, logging attribute, or storage characteristics of index partition <i>partition</i> or its subpartitions.</p> <p><b>Restriction:</b> You cannot specify the <i>physical_attributes_clause</i> for an index on a hash-partitioned table.</p> <p><b>Note:</b> If the index is a local index on a composite-partitioned table, the changes you specify here will override any attributes specified earlier for the subpartitions of index, as well as establish default values of attributes for future subpartitions of that partition. To change the default attributes of the partition without overriding the attributes of subpartitions, use ALTER TABLE ... MODIFY DEFAULT ATTRIBUTES OF PARTITION.</p>
<i>rename_partition/subpartition_clause</i>	<p>renames index partition or subpartition to <i>new_name</i>.</p>
<i>drop_partition_clause</i>	<p>removes a partition and the data in it from a partitioned global index. When you drop a partition of a global index, Oracle marks the index's next partition UNUSABLE. You cannot drop the highest partition of a global index.</p>
<i>split_partition_clause</i>	<p>splits a partition of a global partitioned index into two partitions, adding a new partition to the index.</p> <p>Splitting a partition marked UNUSABLE results in two partitions, both marked UNUSABLE. You must rebuild the partitions before you can use them.</p> <p>Splitting a usable partition results in two partitions populated with index data. Both new partitions are usable.</p>

---

---

	<code>AT (<i>value_list</i>)</code>	specifies the new noninclusive upper bound for <i>split_partition_1</i> . The <i>value_list</i> must evaluate to less than the presplit partition bound for <i>partition_name_old</i> and greater than the partition bound for the next lowest partition (if there is one).
	<code>INTO</code>	describes the two partitions resulting from the split.
<i>partition_</i> <i>description</i>		specifies (optionally) the name and physical attributes of each of two partitions resulting from a split.
<i>modify_</i> <i>subpartition_clause</i>		lets you mark UNUSABLE or allocate or deallocate storage for a subpartition of a local index on a composite-partitioned table. All other attributes of such a subpartition are inherited from partition-level default attributes.

---

## Examples

**Modifying Real Attributes** This statement alters SCOTT's CUSTOMER index so that future data blocks within this index use 5 initial transaction entries and an incremental extent of 100 kilobytes:

```
ALTER INDEX scott.customer
  INITRANS 5
  STORAGE (NEXT 100K);
```

If the SCOTT.CUSTOMER index is partitioned, this statement also alters the default attributes of future partitions of the index. New partitions added in the future will use 5 initial transaction entries and an incremental extent of 100K.

**Dropping an Index Partition** The following statement drops index partition IX\_ANTARTICA:

```
ALTER INDEX sales_area_ix
  DROP PARTITION ix_antarctica;
```

**Modifying Default Attributes** This statement alters the default attributes of local partitioned index SALES\_IX3. New partitions added in the future will use 5 initial transaction entries and an incremental extent of 100K:

```
ALTER INDEX sales_ix3
  MODIFY DEFAULT ATTRIBUTES INITRANS 5 STORAGE ( NEXT 100K );
```

**Marking an Index Unusable** The following statement marks the IDX\_ACCTNO index as UNUSABLE:

```
ALTER INDEX idx_acctno UNUSABLE;
```

**Marking a Partition Unusable** The following statement marks partition `IDX_FEB96` of index `IDX_ACCTNO` as `UNUSABLE`:

```
ALTER INDEX idx_acctno MODIFY PARTITION idx_feb96 UNUSABLE;
```

**Changing MAXEXTENTS** The following statement changes the maximum number of extents for partition `BRIX_NY` and changes the logging attribute:

```
ALTER INDEX branch_ix MODIFY PARTITION brix_ny
  STORAGE( MAXEXTENTS 30 ) LOGGING;
```

**Disabling Parallel Queries** The following statement sets the parallel attributes for index `ARTIST_IX` so that scans on the index will not be parallelized:

```
ALTER INDEX artist_ix NOPARALLEL;
```

**Rebuilding a Partition** The following statement rebuilds partition `P063` in index `ARTIST_IX`. The rebuilding of the index partition will not be logged:

```
ALTER INDEX artist_ix
  REBUILD PARTITION p063 NOLOGGING;
```

**Renaming an Index** The following statement renames an index:

```
ALTER INDEX emp_ix1 RENAME TO employee_ix1;
```

**Renaming an Index Partition** The following statement renames an index partition:

```
ALTER INDEX employee_ix1 RENAME PARTITION emp_ix1_p3
  TO employee_ix1_p3;
```

**Splitting a Partition** The following statement splits partition `PARTNUM_IX_P6` in partitioned index `PARTNUM_IX` into `PARTNUM_IX_P5` and `PARTNUM_IX_P6`:

```
ALTER INDEX partnum_ix
  SPLIT PARTITION partnum_ix_p6 AT ( 5001 )
  INTO ( PARTITION partnum_ix_p5 TABLESPACE ts017 LOGGING,
        PARTITION partnum_ix_p6 TABLESPACE ts004 );
```

The second partition retains the name of the old partition.

**Storing Index Blocks in Reverse Order** The following statement rebuilds index `EMP_IX` so that the bytes of the index block are stored in `REVERSE` order:

```
ALTER INDEX emp_ix REBUILD REVERSE;
```

**Collecting Index Statistics** The following statement collects statistics on the nonpartitioned EMP\_INDX index:

```
ALTER INDEX emp_idx REBUILD COMPUTE STATISTICS;
```

The type of statistics collected depends on the type of index you are rebuilding.

**See Also:** *Oracle8i Concepts*.

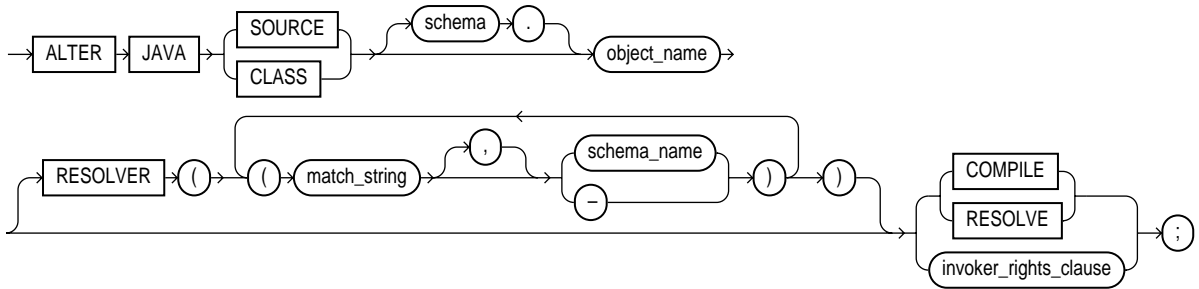
**PARALLEL Example** The following statement causes the index to be rebuilt from the existing index by using parallel execution processes to scan the old and to build the new index:

```
ALTER INDEX emp_idx  
  REBUILD  
  PARALLEL;
```




## ALTER JAVA

### Syntax



### invoker\_rights\_clause::=



### Purpose

To force the resolution of a Java class schema object or compilation of a Java source schema object. (You cannot call the methods of a Java class before all its external references to Java names are associated with other classes.)

**See Also:** *Oracle8i Java Stored Procedures Developer's Guide* for more information on resolving Java classes and compiling Java sources.

### Prerequisites

The Java source or class must be in your own schema, or you must have the ALTER ANY PROCEDURE system privilege. You must also have the EXECUTE object privilege on Java classes.

### Keywords and Parameters

---

JAVA SOURCE    compiles a Java source schema object.

---

JAVA CLASS	resolves a Java class schema object.
<i>object_name</i>	specifies a previously created Java class or source schema object. Use double quotation marks to preserve lower- or mixed-case names.
RESOLVER	specifies how schemas are searched for referenced fully specified Java names, using the mapping pairs specified when the Java class or source was created. <b>See Also:</b> <a href="#">"CREATE JAVA"</a> on page 7-311.
RESOLVE   COMPILE	are synonymous keywords. They specify that Oracle should attempt to resolve the primary Java class schema object. <ul style="list-style-type: none"><li>■ When applied to a class, resolution of referenced names to other class schema objects occurs.</li><li>■ When applied to a source, source compilation occurs.</li></ul>
<i>invoker_rights_clause</i>	specifies whether the methods of the class execute with the privileges and in the schema of the user who defined it or with the privileges and in the schema of <code>CURRENT_USER</code> . For information on how <code>CURRENT_USER</code> is determined, see <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i> .  This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the member functions and procedures of the type. <b>See Also:</b> <i>Oracle8i Java Stored Procedures Developer's Guide</i> .
AUTHID CURRENT_USER	specifies that the methods of the class execute with the privileges of <code>CURRENT_USER</code> . This clause is the default and creates an "invoker-rights class."  This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of <code>CURRENT_USER</code> . External names in all other statements resolve in the schema in which the methods reside.
AUTHID DEFINER	specifies that the methods of the class execute with the privileges of the user who defined it.  This clause also specifies that external names resolve in the schema where the methods reside.

---

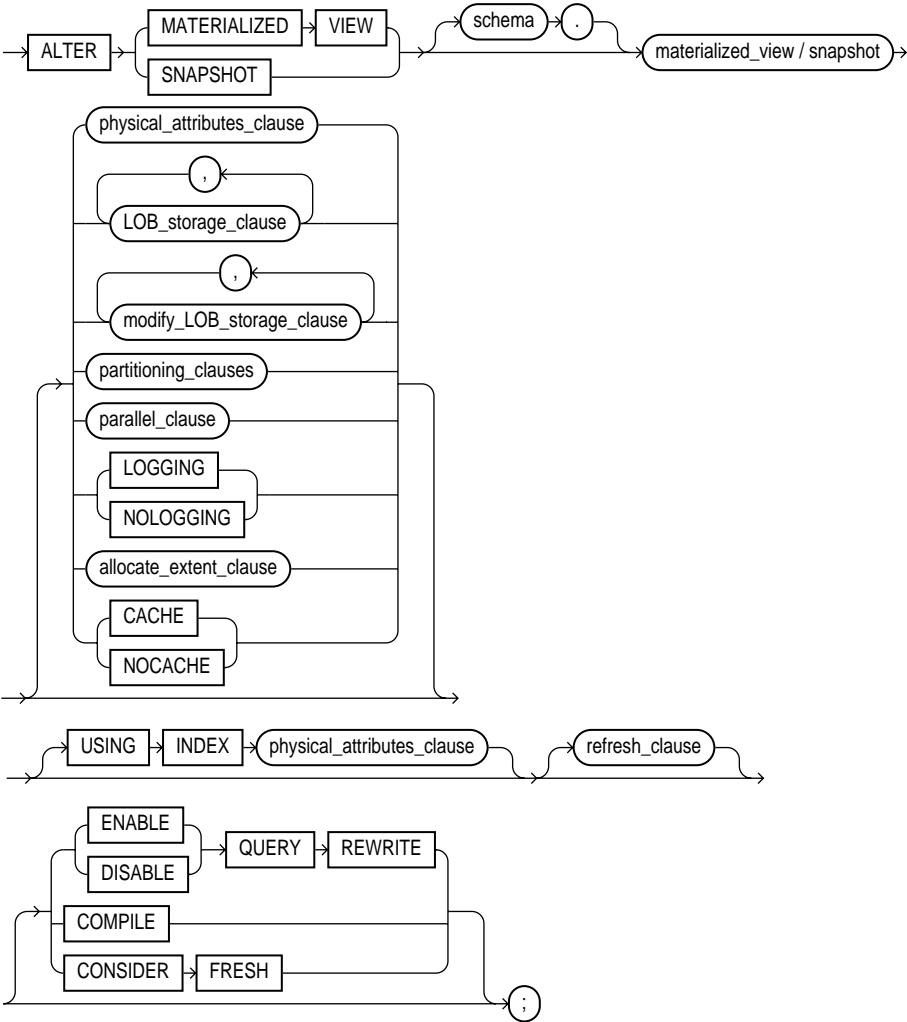
## Example

The following statement forces the resolution of a Java class:

```
ALTER JAVA CLASS "Agent"
  RESOLVER ("/home/java/bin/*" scott)(* public)
  RESOLVE;
```

# ALTER MATERIALIZED VIEW / SNAPSHOT

## Syntax




**LOB\_storage\_clause:** See "ALTER TABLE" on page 7-123.

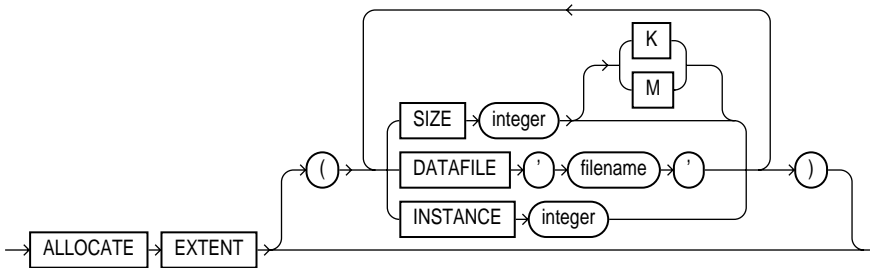
**modify\_LOB\_storage\_clause:** See "ALTER TABLE" on page 7-123.

**partitioning\_clauses:** See "ALTER TABLE" on page 7-123.

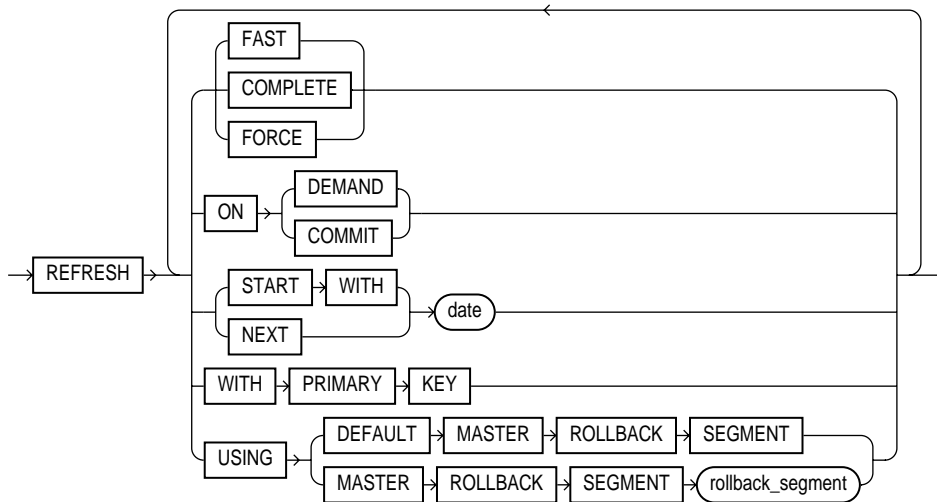
**parallel\_clause::=**

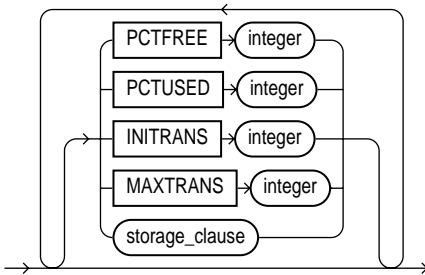


**allocate\_extent\_clause::=**



**refresh\_clause::=**



**physical\_attributes\_clause::=**

**storage\_clause:** See the "storage\_clause" on page 7-605.

**Purpose**

To modify an existing materialized view in one or more of the following ways:

- To change its storage characteristics
- To change its refresh method, mode, or time
- To alter its structure so that it is a different type of materialized view
- To enable or disable query rewrite.

The terms **snapshot** and **materialized view** are synonymous in Oracle documentation. "Materialized view" is used in this reference. Both refer to a database object that contains the results of a query of one or more tables.

The tables in the query are called **master tables** (a replication term) or **detail tables** (a data warehouse term). This reference uses "master tables" for consistency. The databases containing the master tables are called the **master databases**.

**See Also:**

- ["CREATE MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-318 for more information on creating materialized views.
- *Oracle8i Replication* for information on materialized views in a replication environment.
- *Oracle8i Data Warehousing Guide* for information on materialized views in a data warehousing environment.

## Prerequisites

The privileges required to alter a materialized view should be granted directly, as follows:

The materialized view must be in your own schema, or you must have the ALTER ANY SNAPSHOT or ALTER ANY MATERIALIZED VIEW system privilege.

To enable a materialized view for query rewrite:

- If all of the master tables in the materialized view are in your schema, you must have the QUERY REWRITE privilege.
- If any of the master tables are in another schema, you must have the GLOBAL QUERY REWRITE privilege.
- If the materialized view is in another user's schema, both you and the owner of that schema must have the appropriate QUERY REWRITE privilege, as described in the preceding two items. In addition, the owner of the materialized view must have SELECT access to any master tables that the materialized view owner does not own.

**See Also:** *Oracle8i Replication* and *Oracle8i Data Warehousing Guide*.

## Keywords and Parameters

---

<i>schema</i>	is the schema containing the materialized view. If you omit <i>schema</i> , Oracle assumes the materialized view is in your own schema.
<i>materialized view / snapshot</i>	is the name of the materialized view to be altered.
<i>physical_attributes_clause</i>	changes values for the PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters (or, when used in the USING INDEX clause, for the INITRANS and MAXTRANS parameters only) and the storage characteristics for the materialized view. <b>See Also:</b> <ul style="list-style-type: none"><li>■ <a href="#">"ALTER TABLE"</a> on page 7-123 for information on the PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters.</li><li>■ <a href="#">"storage_clause"</a> on page 7-605 for information about storage characteristics.</li></ul>
LOGGING   NOLOGGING	establishes or changes the logging characteristics of the materialized view. <b>See Also:</b> <a href="#">"ALTER TABLE"</a> on page 7-123 for information about logging characteristics.
<i>allocate_extent_clause</i>	explicitly allocates a new extent for the materialized view. <b>See Also:</b> <a href="#">"ALTER TABLE"</a> on page 7-123.

---

**CACHE** | For data that will be accessed frequently, **CACHE** specifies that the blocks retrieved for this table are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables. **NOCACHE** specifies that the blocks are placed at the least recently used end of the LRU list.

**See Also:** "[ALTER TABLE](#)" on page 7-123 for information about specifying **CACHE** or **NOCACHE**.

**LOB\_storage\_clause** specifies the LOB storage characteristics.

**See Also:** "[ALTER TABLE](#)" on page 7-123 for information about specifying the parameters of this clause.

**modify\_LOB\_storage\_clause** modifies the physical attributes of the LOB attribute *lob\_item* or LOB object attribute.

**See Also:** "[ALTER TABLE](#)" on page 7-123 for information about specifying the parameters of this clause.

**partitioning clauses:** The syntax and general functioning of the partitioning clauses is the same as for partitioned tables, as described in "[ALTER TABLE](#)" on page 7-123.

**Restrictions:**

- You cannot use the *LOB\_storage\_clause* or *modify\_LOB\_storage\_clause* when modifying a materialized view.
- If you attempt to drop, truncate, or exchange a materialized view partition, Oracle raises an error.

---

**Note:** If you wish to keep the contents of the materialized view synchronized with those of the master table, Oracle Corporation recommends that you manually perform a complete refresh of all materialized views dependent on the table after dropping or truncating a table partition.

---

**parallel\_clause** changes the default degree of parallelism for the materialized view.

**NOPARALLEL** specifies serial execution.

**PARALLEL** causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances multiplied by the value of the **PARALLEL\_THREADS\_PER\_CPU** initialization parameter.

**PARALLEL integer** specifies the **degree of parallelism**, which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution processes. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify *integer*.

**See Also:** The [Notes](#) to the *parallel\_clause* of "[CREATE TABLE](#)" on page 4-381.

**MODIFY PARTITION UNUSABLE LOCAL INDEXES**

marks **UNUSABLE** all the local index partitions associated with *partition*.

---

**MODIFY PARTITION REBUILD UNUSABLE LOCAL INDEXES**

rebuilds the unusable local index partitions associated with *partition*.

**USING INDEX** changes the value of `INITTRANS`, `MAXTRANS`, and `STORAGE` parameters for the index Oracle uses to maintain the materialized view's data.

**Restriction:** You cannot specify the `PCTUSED` or `PCTFREE` parameters in this clause.

*refresh\_clause* changes the default method and mode and the default times for automatic refreshes. If the contents of a materialized view's master tables are modified, the data in the materialized view must be updated to make the materialized view accurately reflect the data currently in its master table(s). This clause lets you schedule the times and specify the method and mode for Oracle to refresh the materialized view.

---

**Note:** This clause only sets the default refresh options. For instructions on actually implementing the refresh, refer to *Oracle8i Replication* and *Oracle8i Data Warehousing Guide*.

---

**FAST** specifies the incremental refresh method, which performs the refresh according to the changes that have occurred to the master tables. The changes are stored either in the materialized view log associated with the master table (for conventional DML changes) or in the direct loader log (for direct-load INSERTs).

For both conventional DML changes and for direct-path loads, other conditions may restrict the eligibility of a materialized view for fast refresh.

**See Also:**

- *Oracle8i Replication* for restrictions on fast refresh in replication environments
- *Oracle8i Data Warehousing Guide* for restrictions on fast refresh in data warehouse environments

---

**Notes:**

- When you specify `FAST` refresh at create time, Oracle verifies that the materialized view you are creating is eligible for fast refresh. When you change the refresh method to `FAST` in an `ALTER MATERIALIZED VIEW` statement, Oracle does not perform this verification. If the materialized view is not eligible for fast refresh, Oracle will return an error when you attempt to refresh this view.
- Materialized views are not eligible for fast refresh if the defining query contains an analytic function.

**See Also:** "[Analytic Functions](#)" on page 4-7.

---

**COMPLETE** specifies the complete refresh method, which is implemented by executing the materialized view's defining query. If you request a complete refresh, Oracle performs a complete refresh even if a fast refresh is possible.



---

FORCE	specifies that when a refresh occurs, Oracle will perform a fast refresh if one is possible or a complete refresh otherwise.
ON COMMIT	<p>specifies that a fast refresh is to occur whenever Oracle commits a transaction that operates on a master table of the materialized view.</p> <p><b>Restriction:</b> This clause is supported only for materialized join views and single-table materialized aggregate views.</p> <p><b>See Also:</b> <i>Oracle8i Replication</i> and <i>Oracle8i Data Warehousing Guide</i>.</p>
ON DEMAND	<p>specifies that the materialized view will be refreshed on demand by calling one of the three DBMS_MVIEW refresh procedures. If you omit both ON COMMIT and ON DEMAND, ON DEMAND is the default.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ <i>Oracle8i Supplied PL/SQL Packages Reference</i> for information on these procedures.</li> <li>■ <i>Oracle8i Data Warehousing Guide</i> on the types of materialized views you can create by specifying REFRESH ON DEMAND.</li> </ul> <p>If you specify ON COMMIT or ON DEMAND, you cannot also specify START WITH or NEXT.</p>
START WITH	specifies a date expression for the first automatic refresh time.
NEXT	specifies a date expression for calculating the interval between automatic refreshes.
<p>Both the START WITH and NEXT values must evaluate to a time in the future. If you omit the START WITH value, Oracle determines the first automatic refresh time by evaluating the NEXT expression with respect to the creation time of the materialized view. If you specify a START WITH value but omit the NEXT value, Oracle refreshes the materialized view only once. If you omit both the START WITH and NEXT values, or if you omit the <i>refresh_clause</i> entirely, Oracle does not automatically refresh the materialized view.</p>	
WITH PRIMARY KEY	<p>changes a rowid materialized view to a primary key materialized view. Primary key materialized views allow materialized view master tables to be reorganized without affecting the materialized view's ability to continue to fast refresh. The master table must contain an enabled primary key constraint.</p> <p><b>See Also:</b> <i>Oracle8i Replication</i> for detailed information about primary key materialized views.</p>
USING ROLLBACK SEGMENT	<p>changes the remote rollback segment to be used during materialized view refresh, where <i>rollback_segment</i> is the name of the rollback segment to be used.</p> <p><b>See Also:</b> <i>Oracle8i Replication</i> for information on changing the local materialized view rollback segment using the DBMS_REFRESH package.</p>

- **DEFAULT** specifies that Oracle will choose the rollback segment to use. If you specify **DEFAULT**, you cannot specify *rollback\_segment*.
- **MASTER** specifies the remote rollback segment to be used at the remote master for the individual materialized view. (To change the local materialized view rollback segment, use the **DBMS\_REFRESH** package, described in *Oracle8i Replication*.)

The master rollback segment is stored on a per-materialized-view basis and is validated during materialized view creation and refresh. If the materialized view is complex, the master rollback segment, if specified, is ignored.

**QUERY REWRITE** specifies whether the materialized view is eligible to be used for query rewrite.

**ENABLE** enables the materialized view for query rewrite.

**See Also:** *Oracle8i Data Warehousing Guide* for more information on query rewrite.

**Restrictions:**

- If the materialized view is in an invalid or unusable state, it is not eligible for query rewrite in spite of the **ENABLE** mode.
- You cannot enable query rewrite if the materialized view was created totally or in part from a view.
- You can enable query rewrite only if all user-defined functions in the materialized view are **DETERMINISTIC**.

**See Also:** "**CREATE FUNCTION**" on page 7-284.

- You can enable query rewrite only if expressions in the statement are repeatable. For example, you cannot include **CURRENT\_TIME** or **USER**.

**See Also:** *Oracle8i Data Warehousing Guide*.

**DISABLE** specifies that the materialized view is not eligible for use by query rewrite. (If a materialized view is in the invalid state, it is not eligible for use by query rewrite, whether or not it is disabled.) However, a disabled materialized view can be refreshed.

**COMPILE** explicitly revalidates a materialized view. If an object upon which the materialized view depends is dropped or altered, the materialized view remains accessible, but it is invalid for query rewrite. You can use this clause to explicitly revalidate the materialized view to make it eligible for query rewrite.

If the materialized view fails to revalidate, it cannot be refreshed or used for query rewrite.

CONSIDER  
FRESH

directs Oracle to consider the materialized view fresh and therefore eligible for query rewrite in the `TRUSTED` or `STALE_TOLERATED` modes. Because Oracle cannot guarantee the freshness of the materialized view, query rewrite in `ENFORCED` mode is not supported. This clause also sets the staleness state of the materialized view to `UNKNOWN`. The staleness state is displayed in the `STALENESS` column of the `ALL_MVIEWS`, `DBA_MVIEWS`, and `USER_MVIEWS` data dictionary views.

This clause is useful after performing partition maintenance operations against the master table. Such operations would otherwise render the materialized view ineligible for fast refresh, and eligible for query rewrite only in `STALE_TOLERATED` mode.

**Note:** A materialized view is stale if changes have been made to the contents of any of its master tables. This clause directs Oracle to assume that the materialized view is fresh and that no such changes have been made. Therefore, actual updates to those tables pending refresh are purged with respect to the materialized view.

---

**See Also:** *Oracle8i Data Warehousing Guide* for more information on query rewrite and the implications of performing partition maintenance operations on master tables.

---

## Examples

**Automatic Refresh Example** The following statement changes the default refresh method for the `HQ_EMP` materialized view to `FAST`:

```
CREATE MATERIALIZED VIEW hq_emp
  REFRESH COMPLETE
  START WITH SYSDATE NEXT SYSDATE +1/4096
  AS SELECT * FROM hq_emp;

ALTER MATERIALIZED VIEW hq_emp
  REFRESH FAST;
```

The next automatic refresh of the materialized view will be a fast refresh provided it is a simple materialized view and its master table has a materialized view log that was created before the materialized view was created or last refreshed.

Because the `REFRESH` clause does not specify `START WITH` or `NEXT` values, the refresh intervals established by the `REFRESH` clause when the `HQ_EMP` materialized view was created or last altered are still used.

**NEXT Example** The following statement stores a new interval between automatic refreshes for the `BRANCH_EMP` materialized view:

```
ALTER MATERIALIZED VIEW branch_emp
  REFRESH NEXT SYSDATE+7;
```

Because the `REFRESH` clause does not specify a `START WITH` value, the next automatic refresh occurs at the time established by the `START WITH` and `NEXT` values specified when the `BRANCH_EMP` materialized view was created or last altered.

At the time of the next automatic refresh, Oracle refreshes the materialized view, evaluates the `NEXT` expression `SYSDATE+7` to determine the next automatic refresh time, and continues to refresh the materialized view automatically once a week.

Because the `REFRESH` clause does not explicitly specify a refresh method, Oracle continues to use the refresh method specified by the `REFRESH` clause of the `CREATE MATERIALIZED VIEW` or most recent `ALTER MATERIALIZED VIEW` statement.

**Complete Refresh Example** The following statement specifies a new refresh method, a new next refresh time, and a new interval between automatic refreshes of the `SF_EMP` materialized view:

```
ALTER MATERIALIZED VIEW sf_emp
  REFRESH COMPLETE
  START WITH TRUNC(SYSDATE+1) + 9/24
  NEXT SYSDATE+7;
```

The `START WITH` value establishes the next automatic refresh for the materialized view to be 9:00 a.m. tomorrow. At that point, Oracle performs a complete refresh of the materialized view, evaluates the `NEXT` expression, and subsequently refreshes the materialized view every week.

**Enabling Query Rewrite Example** The following statement enables query rewrite on the materialized view `MV1` and implicitly revalidates it.

```
ALTER MATERIALIZED VIEW mv1
  ENABLE QUERY REWRITE;
```

**Rollback Segment Examples** The following statement changes the remote master rollback segment used during materialized view refresh to `MASTER_SEG`:

```
ALTER MATERIALIZED VIEW inventory
  REFRESH USING MASTER ROLLBACK SEGMENT master_seg;
```

The following statement changes the remote master rollback segment used during materialized view refresh to one chosen by Oracle:

```
ALTER MATERIALIZED VIEW sales
  REFRESH USING DEFAULT MASTER ROLLBACK SEGMENT;
```

**Primary Key Example** The following statement changes a rowid materialized view to a primary key materialized view:

```
ALTER MATERIALIZED VIEW emp_rs  
    REFRESH WITH PRIMARY KEY;
```

**COMPILE Example** The following statement revalidates the materialized view STORE\_MV:

```
ALTER MATERIALIZED VIEW store_mv COMPILE;
```

**Modifying Refresh Method Example** The following statement changes the refresh method of materialized view STORE\_MV to FAST;

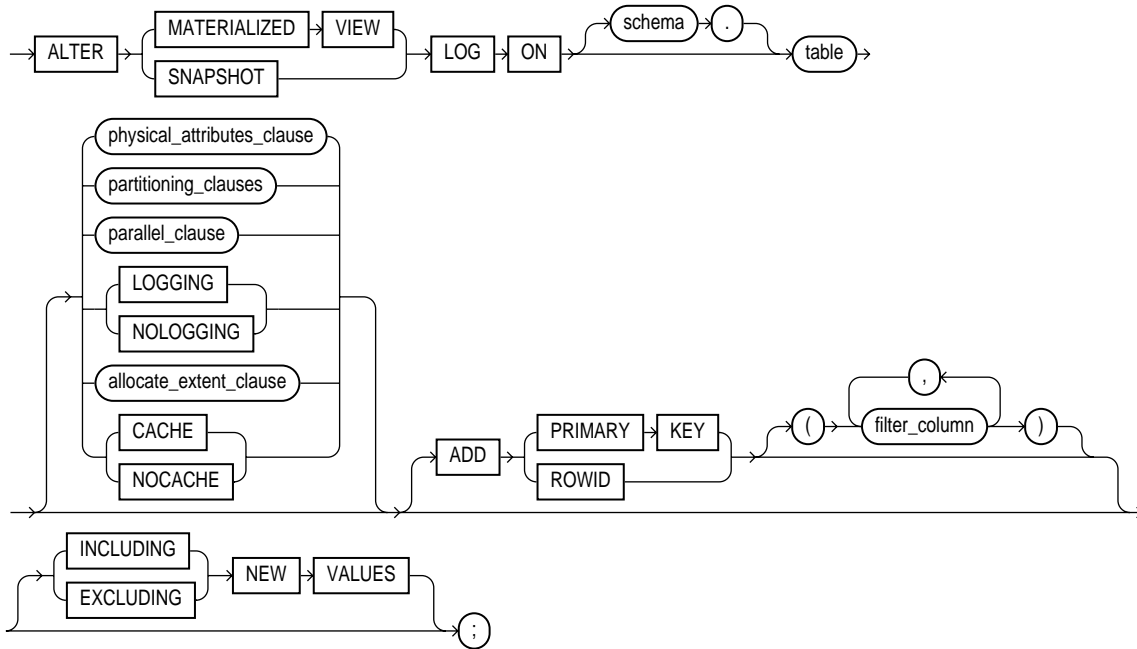
```
ALTER MATERIALIZED VIEW store_mv REFRESH FAST;
```

**CONSIDER FRESH Example** The following statement instructs Oracle that materialized view MV1 should be considered fresh. This statement allows MV1 to be eligible for query rewrite in TRUSTED mode even after you have performed partition maintenance operations on the master tables of MV1:

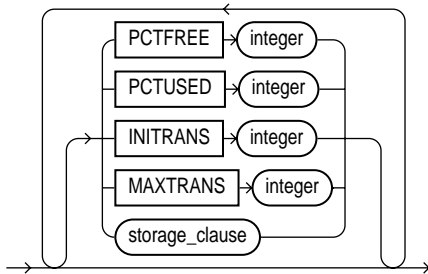
```
ALTER MATERIALIZED VIEW mv1 CONSIDER FRESH;
```

## ALTER MATERIALIZED VIEW LOG / SNAPSHOT LOG

### Syntax



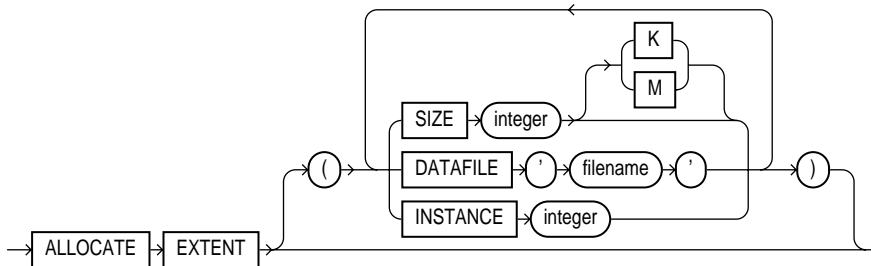
### physical\_attributes\_clause::=




**storage\_clause:** See "[storage\\_clause](#)" on page 7-605.

**partitioning\_clauses:** See "ALTER TABLE" on page 7-123.

**allocate\_extent\_clause::=**



**parallel\_clause::=**



## Purpose

To alter the storage characteristics, refresh mode or time, or type of an existing materialized view log. A **materialized view log** is a table associated with the master table of a materialized view.

The terms **snapshot** and **materialized view** are synonymous. Both refer to a table that contains the results of a query of one or more tables, each of which may be located on the same or on a remote database.

### See Also:

- ["ALTER MATERIALIZED VIEW / SNAPSHOT"](#) on page 2-47 for more information on materialized views, including refreshing them.
- ["CREATE MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-318 for a description of the various types of materialized views

## Prerequisites

Only the owner of the master table or a user with the `SELECT` privilege for the master table can alter a materialized view log.

**See Also:** *Oracle8i Replication* for detailed information about the prerequisites for `ALTER SNAPSHOT LOG`.

## Keywords and Parameters

---

<i>schema</i>	is the schema containing the master table. If you omit <i>schema</i> , Oracle assumes the materialized view log is in your own schema.
<i>table</i>	is the name of the master table associated with the materialized view log to be altered.
<i>physical_attributes_clause</i>	changes the value of <code>PCTFREE</code> , <code>PCTUSED</code> , <code>INITRANS</code> , and <code>MAXTRANS</code> parameters for the table, the partition, the overflow data segment, or the default characteristics of a partitioned table. For a description of these parameters, see " <a href="#">CREATE TABLE</a> " on page 4-381. See also the " <a href="#">Storage Example</a> " on page 2-61.
<i>partitioning_clauses</i>	<p>The syntax and general functioning of the partitioning clauses is the same as for the <code>ALTER TABLE</code> statement; see "<a href="#">ALTER TABLE</a>" on page 7-123.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot use the <i>LOB_storage_clause</i> or <i>modify_LOB_storage_clause</i> when modifying a materialized view log.</li><li>■ If you attempt to drop, truncate, or exchange a materialized view log partition, Oracle raises an error.</li></ul>
<i>parallel_clause</i>	<p>specifies whether parallel operations will be supported for the materialized view log. For additional information, see the <a href="#">Notes</a> to the <i>parallel_clause</i> of "<a href="#">CREATE TABLE</a>" on page 4-381.</p> <p><code>NOPARALLEL</code> specifies serial execution.</p> <p><code>PARALLEL</code> causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the <code>PARALLEL_THREADS_PER_CPU</code> initialization parameter.</p> <p><code>PARALLEL integer</code> specifies the <b>degree of parallelism</b>, which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution processes. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i>.</p>
<code>LOGGING</code>   <code>NOLOGGING</code>	specifies the logging attribute. For information about specifying this attribute, see " <a href="#">ALTER TABLE</a> " on page 7-123.



---

<i>allocate_extent_clause</i>	explicitly allocates a new extent for the materialized view log. <b>See Also:</b> "ALTER TABLE" on page 7-123.
CACHE   NOCACHE	For data that will be accessed frequently, CACHE specifies that the blocks retrieved for this log are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables. NOCACHE specifies that the blocks are placed at the least recently used end of the LRU list. <b>See Also:</b> "ALTER TABLE" on page 7-123 for information about specifying CACHE or NOCACHE.
ADD	augments the materialized view log so that it records the primary key values or rowid values when rows in the materialized view master table are updated. This clause can also be used to record additional filter columns.  To stop recording any of this information, you must first drop the materialized view log and then re-create it. Dropping the materialized view log and then re-creating it forces each of the existing materialized views that depend on the master table to complete refresh on its next refresh.
PRIMARY KEY	specifies that the primary-key values of all rows that are updated should be recorded in the materialized view log.
ROWID	specifies that the rowid values of all rows that are updated should be recorded in the materialized view log.
<i>filter_column(s)</i>	specifies that the values of these columns should be recorded in the materialized view log for all rows that are updated. Filter columns are non-primary-key columns referenced by materialized views. For information about filter columns, see <i>Oracle8i Replication</i> .
NEW VALUES	specifies whether Oracle saves both old and new values in the materialized view log.
INCLUDING	saves both new and old values in the log. If this log is for a table on which you have a single-table materialized aggregate view, and if you want the materialized view to be eligible for fast refresh, you must specify INCLUDING.
EXCLUDING	disables the recording of new values in the log. You can use this clause to avoid the overhead of recording new values. However, do not use this clause if you have a fast-refreshable single-table materialized aggregate view defined on this table.

---

## Examples

**Storage Example** The following statement changes the MAXEXTENTS value of a materialized view log:


```
ALTER MATERIALIZED VIEW LOG ON dept
  STORAGE MAXEXTENTS 50;
```

**PRIMARY KEY Example** The following statement alters an existing rowid materialized view log to also record primary key information:

```
ALTER MATERIALIZED VIEW LOG ON sales  
  ADD PRIMARY KEY;
```

## ALTER OUTLINE

### Syntax



### Purpose

To rename a stored outline, reassign it to a different category, or regenerate it by compiling the outline's SQL statement and replacing the old outline data with the outline created under current conditions.

**See Also:** ["CREATE OUTLINE"](#) on page 7-342 and *Oracle8i Designing and Tuning for Performance* for more information on outlines.

### Prerequisites

To modify an outline, you must have the `ALTER ANY OUTLINE` system privilege.

### Keywords and Parameters

<i>outline</i>	is the name of the outline to be modified.
REBUILD	regenerates the execution plan for <i>outline</i> using current conditions.
RENAME TO <i>new_outline_name</i>	specifies an outline name to replace <i>outline</i> .
CHANGE CATEGORY TO <i>new_category_name</i>	specifies the name of the category into which the <i>outline</i> will be moved.


### Example

The following statement regenerates a stored outline called `SALARIES` by compiling the outline's text and replacing the old outline data with the outline created under current conditions.

```
ALTER OUTLINE salaries REBUILD;
```

## ALTER PACKAGE

### Syntax



### Purpose

To explicitly recompile either a package specification, body, or both. Explicit recompilation eliminates the need for implicit run-time recompilation and prevents associated run-time compilation errors and performance overhead.

Because all objects in a package are stored as a unit, the `ALTER PACKAGE` statement recompiles all package objects together. You cannot use the `ALTER PROCEDURE` statement or `ALTER FUNCTION` statement to recompile individually a procedure or function that is part of a package.

---



---

**Note:** This statement does not change the declaration or definition of an existing package. To redeclare or redefine a package, use the "[CREATE PACKAGE](#)" or the "[CREATE PACKAGE BODY](#)" statement with the `OR REPLACE` clause.

---



---

### Prerequisites

The package must be in your own schema or you must have `ALTER ANY PROCEDURE` system privilege.

### Keywords and Parameters

<i>schema</i>	is the schema containing the package. If you omit <i>schema</i> , Oracle assumes the package is in your own schema.
<i>package</i>	is the name of the package to be recompiled.

---

COMPILE	<p>recompiles the package specification or body. The <code>COMPILE</code> keyword is required.</p> <p>If recompiling the package results in compilation errors, Oracle returns an error and the body remains invalid. You can see the associated compiler error messages with the <code>SQL*Plus</code> command <code>SHOW ERRORS</code>.</p>
SPECIFICATION	<p>recompiles only the package specification, regardless of whether it is invalid. You might want to recompile a package specification to check for compilation errors after modifying the specification.</p> <p>When you recompile a package specification, Oracle invalidates any local objects that depend on the specification, such as procedures that call procedures or functions in the package. The body of a package also depends on its specification. If you subsequently reference one of these dependent objects without first explicitly recompiling it, Oracle recompiles it implicitly at run time.</p>
BODY	<p>recompiles only the package body regardless of whether it is invalid. You might want to recompile a package body after modifying it. Recompiling a package body does not invalidate objects that depend upon the package specification.</p> <p>When you recompile a package body, Oracle first recompiles the objects on which the body depends, if any of those objects are invalid. If Oracle recompiles the body successfully, the body becomes valid.</p>
PACKAGE	<p>recompiles both the package specification and the package body if one exists, regardless of whether they are invalid. This is the default. The recompilation of the package specification and body lead to the invalidation and recompilation as described above for <code>SPECIFICATION</code> and <code>BODY</code>.</p> <p>For information on how Oracle maintains dependencies among schema objects, including remote objects, see <i>Oracle8i Concepts</i>.</p>
DEBUG	<p>instructs the PL/SQL compiler to generate and store the code for use by the PL/SQL debugger.</p> <p>For information on debugging packages, see <i>Oracle8i Supplied PL/SQL Packages Reference</i>.</p>

---

## Examples

This statement explicitly recompiles the specification and body of the `ACCOUNTING` package in the schema `BLAIR`:

```
ALTER PACKAGE blair.accounting
  COMPILE PACKAGE;
```

If Oracle encounters no compilation errors while recompiling the `ACCOUNTING` specification and body, `ACCOUNTING` becomes valid. `BLAIR` can subsequently call or reference all package objects declared in the specification of `ACCOUNTING` without run-time recompilation. If recompiling `ACCOUNTING` results in compilation errors, Oracle returns an error and `ACCOUNTING` remains invalid.

Oracle also invalidates all objects that depend upon `ACCOUNTING`. If you subsequently reference one of these objects without explicitly recompiling it first, Oracle recompiles it implicitly at run time.

To recompile the body of the `ACCOUNTING` package in the schema `BLAIR`, issue the following statement:

```
ALTER PACKAGE blair.accounting  
  COMPILE BODY;
```


If Oracle encounters no compilation errors while recompiling the package body, the body becomes valid. `BLAIR` can subsequently call or reference all package objects declared in the specification of `ACCOUNTING` without run-time recompilation. If recompiling the body results in compilation errors, Oracle returns an error message and the body remains invalid.

Because this statement recompiles the body and not the specification of `ACCOUNTING`, Oracle does not invalidate dependent objects.

---

## ALTER PROCEDURE

### Syntax



### Purpose

To explicitly recompile a stand-alone stored procedure. Explicit recompilation eliminates the need for implicit run-time recompilation and prevents associated run-time compilation errors and performance overhead.

To recompile a procedure that is part of a package, recompile the entire package using the `ALTER PACKAGE` statement (see "[ALTER PACKAGE](#)" on page 2-64).

---

**Note:** This statement does not change the declaration or definition of an existing procedure. To redeclare or redefine a procedure, use the `CREATE PROCEDURE` statement with the `OR REPLACE` clause (see "[CREATE PROCEDURE](#)" on page 7-353)

---

The `ALTER PROCEDURE` statement is quite similar to the `ALTER FUNCTION` statement (see "[ALTER FUNCTION](#)" on page 2-28).

### Prerequisites

The procedure must be in your own schema or you must have `ALTER ANY PROCEDURE` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the procedure. If you omit <i>schema</i> , Oracle assumes the procedure is in your own schema.
<i>procedure</i>	is the name of the procedure to be recompiled.

COMPILE	<p>causes Oracle to recompile the procedure. The <code>COMPILE</code> keyword is required. Oracle recompiles the procedure regardless of whether it is valid or invalid.</p> <ul style="list-style-type: none"><li>■ Oracle first recompiles objects upon which the procedure depends, if any of those objects are invalid.</li><li>■ Oracle also invalidates any local objects that depend upon the procedure, such as procedures that call the recompiled procedure or package bodies that define procedures that call the recompiled procedure.</li><li>■ If Oracle recompiles the procedure successfully, the procedure becomes valid. If recompiling the procedure results in compilation errors, then Oracle returns an error and the procedure remains invalid. You can see the associated compiler error messages with the SQL*Plus command <code>SHOW ERRORS</code>.</li></ul> <p>For information on how Oracle maintains dependencies among schema objects, including remote objects, see <i>Oracle8i Concepts</i>.</p>
DEBUG	<p>instructs the PL/SQL compiler to generate and store the code for use by the PL/SQL debugger.</p> <p>For information on debugging procedures, see <i>Oracle8i Application Developer's Guide - Fundamentals</i>.</p>

---

### Example

To explicitly recompile the procedure `CLOSE_ACCT` owned by the user `HENRY`, issue the following statement:

```
ALTER PROCEDURE henry.close_acct
  COMPILE;
```


If Oracle encounters no compilation errors while recompiling `CLOSE_ACCT`, `CLOSE_ACCT` becomes valid. Oracle can subsequently execute it without recompiling it at run time. If recompiling `CLOSE_ACCT` results in compilation errors, Oracle returns an error and `CLOSE_ACCT` remains invalid.

Oracle also invalidates all dependent objects. These objects include any procedures, functions, and package bodies that call `CLOSE_ACCT`. If you subsequently reference one of these objects without first explicitly recompiling it, Oracle recompiles it implicitly at run time.

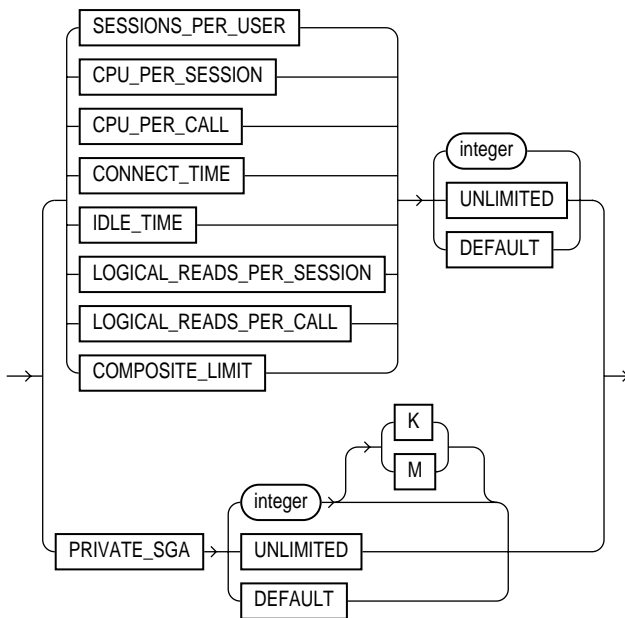


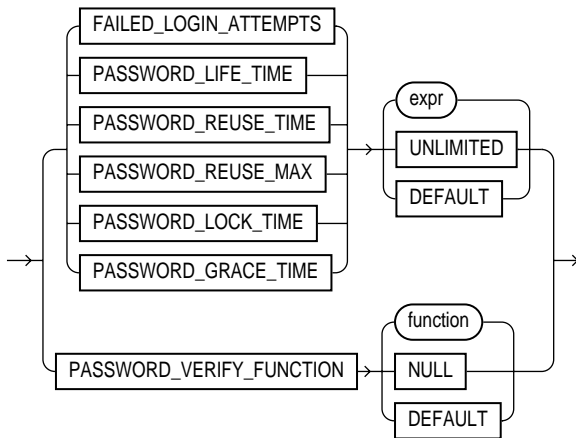
# ALTER PROFILE

## Syntax



### resource\_parameters::=



**password\_parameters::=****Purpose**

To add, modify, or remove a resource limit or password management parameter in a profile.

Changes made to a profile with an `ALTER PROFILE` statement affect users only in their subsequent sessions, not in their current sessions.

For information on creating a profile, see "[CREATE PROFILE](#)" on page 7-359.

**Prerequisites**

You must have `ALTER PROFILE` system privilege to change profile resource limits. To modify password limits and protection, you must have `ALTER PROFILE` and `ALTER USER` system privileges.

**Keywords and Parameters**

The keywords and parameters in the `ALTER PROFILE` statement all have the same meaning as in the `CREATE PROFILE` statement. See "[CREATE PROFILE](#)" on page 7-359.

---



---

**Note:** You cannot remove a limit from the `DEFAULT` profile.

---



---

## Examples

**Making a Password Unavailable** The following statement makes a password unavailable for reuse for 90 days:

```
ALTER PROFILE prof
  LIMIT PASSWORD_REUSE_TIME 90
  PASSWORD_REUSE_MAX UNLIMITED;
```

**Setting Default Values** The following statement defaults the `PASSWORD_REUSE_TIME` value to its defined value in the `DEFAULT` profile:

```
ALTER PROFILE prof
  LIMIT PASSWORD_REUSE_TIME DEFAULT
  PASSWORD_REUSE_MAX UNLIMITED;
```

**Limiting Login Attempts and Password Lock Time** The following statement alters profile `PROF` with `FAILED_LOGIN_ATTEMPTS` set to 5 and `PASSWORD_LOCK_TIME` set to 1:

```
ALTER PROFILE prof LIMIT
  FAILED_LOGIN_ATTEMPTS 5
  PASSWORD_LOCK_TIME 1;
```

This statement causes `PROF`'s account to become locked for 1 day after 5 unsuccessful login attempts.

**Changing Password Lifetime and Grace Period** The following statement modifies profile `PROF`'s `PASSWORD_LIFE_TIME` to 60 days and `PASSWORD_GRACE_TIME` to 10 days:

```
ALTER PROFILE prof LIMIT
  PASSWORD_LIFE_TIME 60
  PASSWORD_GRACE_TIME 10;
```

**Limiting Concurrent Sessions** This statement defines a new limit of 5 concurrent sessions for the `ENGINEER` profile:

```
ALTER PROFILE engineer LIMIT SESSIONS_PER_USER 5;
```

If the `ENGINEER` profile does not currently define a limit for `SESSIONS_PER_USER`, the above statement adds the limit of 5 to the profile. If the profile already defines a limit, the above statement redefines it to 5. Any user assigned the `ENGINEER` profile is subsequently limited to 5 concurrent sessions.

**Removing Limits** This statement removes the `IDLE_TIME` limit from the `ENGINEER` profile:

```
ALTER PROFILE engineer LIMIT IDLE_TIME DEFAULT;
```

Any user assigned the `ENGINEER` profile is subject in their subsequent sessions to the `IDLE_TIME` limit defined in the `DEFAULT` profile.

**Limiting Idle Time** This statement defines a limit of 2 minutes of idle time for the `DEFAULT` profile:

```
ALTER PROFILE default LIMIT IDLE_TIME 2;
```

This `IDLE_TIME` limit applies to these users:

- Users who are not explicitly assigned any profile
- Users who are explicitly assigned a profile that does not define an `IDLE_TIME` limit


This statement defines unlimited idle time for the `ENGINEER` profile:

```
ALTER PROFILE engineer LIMIT IDLE_TIME UNLIMITED;
```

Any user assigned the `ENGINEER` profile is subsequently permitted unlimited idle time.

## ALTER RESOURCE COST

### Syntax



### Purpose

To specify or change the formula by which Oracle calculates the total resource cost used in a session. The weight that you assign to each resource determines how much the use of that resource contributes to the total resource cost. If you do not assign a weight to a resource, the weight defaults to 0 and use of the resource subsequently does not contribute to the cost. The weights you assign apply to all subsequent sessions in the database.

Oracle calculates the total resource cost by first multiplying the amount of each resource used in the session by the resource's weight, and then summing the products for all four resources. For any session, this cost is limited by the value of the `COMPOSITE_LIMIT` parameter in the user's profile. Both the products and the total cost are expressed in units called **service units**.

Although Oracle monitors the use of other resources, only the four resources shown in the syntax can contribute to the total resource cost for a session. For information on all resources, see ["CREATE PROFILE"](#) on page 7-359.

Once you have specified a formula for the total resource cost, you can limit this cost for a session with the `COMPOSITE_LIMIT` parameter of the `CREATE PROFILE` statement. If a session's cost exceeds the limit, Oracle aborts the session and returns an error. For information on establishing resource limits, see ["CREATE PROFILE"](#) on page 7-359. If you use the `ALTER RESOURCE COST` statement to change the weight assigned to each resource, Oracle uses these new weights to calculate the total resource cost for all current and subsequent sessions.

### Prerequisites

You must have `ALTER RESOURCE COST` system privilege.

## Keywords and Parameters

---

CPU_PER_SESSION	is the amount of CPU time used by a session measured in hundredth of seconds.
CONNECT_TIME	is the elapsed time of a session measured in minutes.
LOGICAL_READS_PER_SESSION	is the number of data blocks read during a session, including blocks read from both memory and disk.
PRIVATE_SGA	is the number of bytes of private space in the system global area (SGA) used by a session. This limit applies only if you are using the multi-threaded server architecture and allocating private space in the SGA for your session.
<i>integer</i>	is the weight of each resource.

---

## Example

The following statement assigns weights to the resources CPU\_PER\_SESSION and CONNECT\_TIME:

```
ALTER RESOURCE COST
  CPU_PER_SESSION 100
  CONNECT_TIME      1;
```

The weights establish this cost formula for a session:

$$\text{cost} = (100 * \text{CPU\_PER\_SESSION}) + (1 * \text{CONNECT\_TIME})$$

where the values of CPU\_PER\_SESSION and CONNECT\_TIME are either values in the DEFAULT profile or in the profile of the user of the session.

Because the above statement assigns no weight to the resources LOGICAL\_READS\_PER\_SESSION and PRIVATE\_SGA, these resources do not appear in the formula.

If a user is assigned a profile with a COMPOSITE\_LIMIT value of 500, a session exceeds this limit whenever COST exceeds 500. For example, a session using 0.04 seconds of CPU time and 101 minutes of elapsed time exceeds the limit. A session 0.0301 seconds of CPU time and 200 minutes of elapsed time also exceeds the limit.

You can subsequently change the weights with another ALTER RESOURCE statement:

```
ALTER RESOURCE COST
  LOGICAL_READS_PER_SESSION 2
  CONNECT_TIME 0;
```

These new weights establish a new cost formula:

```
cost = (100 * CPU_PER_SESSION) + (2 * LOGICAL_READ_PER_SECOND)
```

where the values of `CPU_PER_SESSION` and `LOGICAL_READS_PER_SECOND` are either the values in the `DEFAULT` profile or in the profile of the user of this session.


This `ALTER RESOURCE COST` statement changes the formula in these ways:

- The statement omits a weight for the `CPU_PER_SESSION` resource and the resource was already assigned a weight, so the resource remains in the formula with its original weight.
- The statement assigns a weight to the `LOGICAL_READS_PER_SESSION` resource, so this resource now appears in the formula.
- The statement assigns a weight of 0 to the `CONNECT_TIME` resource, so this resource no longer appears in the formula.
- The statement omits a weight for the `PRIVATE_SGA` resource and the resource was not already assigned a weight, so the resource still does not appear in the formula.

---

## ALTER ROLE

### Syntax



### Purpose

To change the authorization needed to enable a role. For information on creating a role, see ["CREATE ROLE"](#) on page 4-365. For information on enabling or disabling a role for your session, see ["SET ROLE"](#) on page 7-600.

### Prerequisites

You must either have been granted the role with the `ADMIN OPTION` or have `ALTER ANY ROLE` system privilege.

Before you alter a role to `IDENTIFIED GLOBALLY`, you must:

- Revoke all grants of roles identified externally to the role and
- Revoke the grant of the role from all users, roles, and `PUBLIC`.

The one exception to this rule is that you should not revoke the role from the user who is currently altering the role.

### Keywords and Parameters

The keywords and parameters in the `ALTER ROLE` statement all have the same meaning as in the `CREATE ROLE` statement. See ["CREATE ROLE"](#) on page 4-365.



---

---

**Note:** If you have the ALTER ANY ROLE system privilege and you change a role that is IDENTIFIED GLOBALLY to IDENTIFIED BY password, IDENTIFIED EXTERNALLY, or NOT IDENTIFIED, then Oracle grants you the altered role with the ADMIN OPTION, as it would have if you had created the role identified nonglobally.

---

---

## Examples

The following statement changes the role ANALYST to IDENTIFIED GLOBALLY:

```
ALTER ROLE analyst IDENTIFIED GLOBALLY;
```

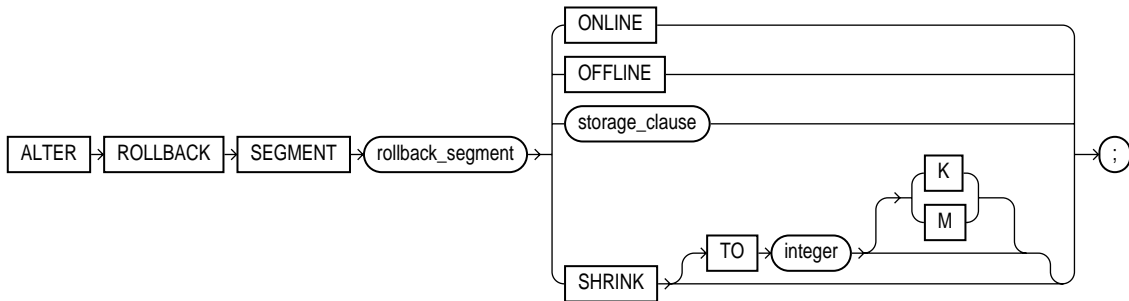
This statement changes the password on the TELLER role to LETTER:

```
ALTER ROLE teller  
    IDENTIFIED BY letter;
```

Users granted the TELLER role must subsequently enter the new password "letter" to enable the role.

## ALTER ROLLBACK SEGMENT

### Syntax



**storage\_clause:** See ["storage\\_clause"](#) on page 7-605.

### Purpose

To bring a rollback segment online or offline, to change its storage characteristics, or to shrink it to an optimal or specified size.

For information on creating a rollback segment, see ["CREATE ROLLBACK SEGMENT"](#) on page 4-367.

### Prerequisites

You must have ALTER ROLLBACK SEGMENT system privilege.

### Keywords and Parameters

---

<i>rollback_segment</i>	specifies the name of an existing rollback segment.
ONLINE	brings the rollback segment online. When you create a rollback segment, it is initially offline and not available for transactions. This clause brings the rollback segment online, making it available for transactions by your instance. You can also bring a rollback segment online when you start your instance with the initialization parameter ROLLBACK_SEGMENTS.

---

OFFLINE	<p>takes the rollback segment offline.</p> <ul style="list-style-type: none"> <li>■ If the rollback segment does not contain any information needed to roll back an active transaction, Oracle takes it offline immediately.</li> <li>■ If the rollback segment does contain information for active transactions, Oracle makes the rollback segment unavailable for future transactions and takes it offline after all the active transactions are committed or rolled back.</li> </ul> <p>Once the rollback segment is offline, it can be brought online by any instance.</p> <p>To see whether a rollback segment is online or offline, query the data dictionary view <code>DBA_ROLLBACK_SEGS</code>. Online rollback segments have a <code>STATUS</code> value of <code>IN_USE</code>. Offline rollback segments have a <code>STATUS</code> value of <code>AVAILABLE</code>.</p> <p><b>Restriction:</b> You cannot take the <code>SYSTEM</code> rollback segment offline.</p> <p><b>See Also:</b> <i>Oracle8i Administrator's Guide</i> for more information on making rollback segments available and unavailable.</p>
<i>storage_clause</i>	<p>changes the rollback segment's storage characteristics. See the "<a href="#">storage_clause</a>" on page 7-605 for syntax and additional information.</p> <p><b>Restriction:</b> You cannot change the values of the <code>INITIAL</code> and <code>MINEXTENTS</code> for an existing rollback segment.</p>
SHRINK	<p>attempts to shrink the rollback segment to an optimal or specified size. The success and amount of shrinkage depend on the available free space in the rollback segment and how active transactions are holding space in the rollback segment.</p> <p>The value of <i>integer</i> is in bytes, unless you specify <code>K</code> or <code>M</code> for kilobytes or megabytes.</p> <p>If you do not specify <code>TO integer</code>, then the size defaults to the <code>OPTIMAL</code> value of the <i>storage_clause</i> of the <code>CREATE ROLLBACK SEGMENT</code> statement that created the rollback segment. If <code>OPTIMAL</code> was not specified, then the size defaults to the <code>MINEXTENTS</code> value of the <i>storage_clause</i> of the <code>CREATE ROLLBACK SEGMENT</code> statement.</p> <p>Regardless of whether you specify <code>TO integer</code>:</p> <ul style="list-style-type: none"> <li>■ The value to which Oracle shrinks the rollback segment is valid for the execution of the statement. Thereafter, the size reverts to the <code>OPTIMAL</code> value of the <code>CREATE ROLLBACK SEGMENT</code> statement.</li> <li>■ The rollback segment cannot shrink to less than two extents.</li> </ul> <p>To determine the actual size of a rollback segment after attempting to shrink it, query the <code>BYTES</code>, <code>BLOCKS</code>, and <code>EXTENTS</code> columns of the <code>DBA_SEGMENTS</code> view.</p> <p><b>Restriction:</b> For a parallel server, you can shrink only rollback segments that are online to your instance.</p>

---

## Examples

This statement brings the rollback segment `RSONE` online:

```
ALTER ROLLBACK SEGMENT rstone ONLINE;
```

**This statement changes the STORAGE parameters for RSONE:**

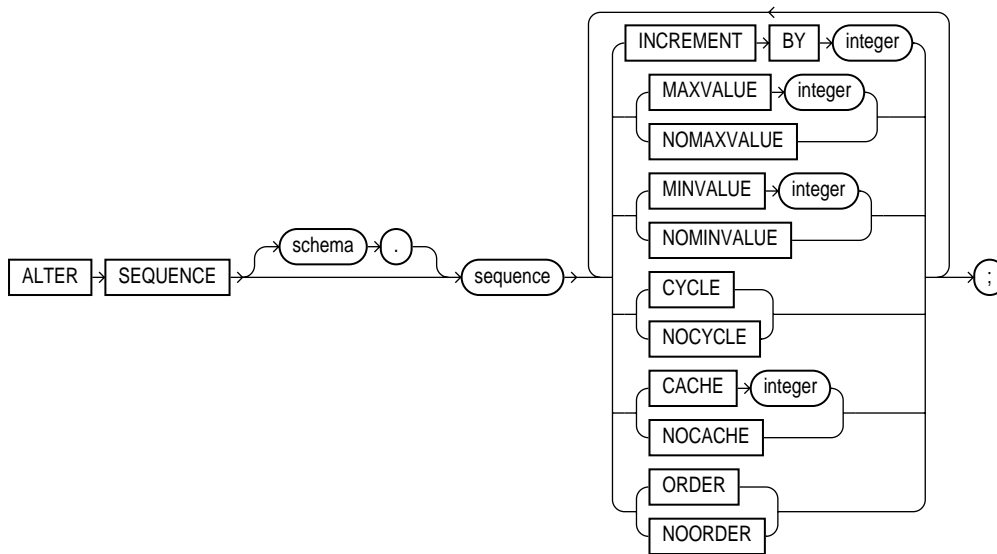
```
ALTER ROLLBACK SEGMENT rsone
    STORAGE (NEXT 1000 MAXEXTENTS 20);
```

**This statement attempts to resize a rollback segment to 100 megabytes:**

```
ALTER ROLLBACK SEGMENT rsone
    SHRINK TO 100 M;
```

## ALTER SEQUENCE

### Syntax



### Purpose

To change the increment, minimum and maximum values, cached numbers, and behavior of an existing sequence. This statement affects only future sequence numbers. For additional information on sequences, see ["CREATE SEQUENCE"](#) on page 4-371.

### Prerequisites

The sequence must be in your own schema, or you must have the ALTER object privilege on the sequence, or you must have the ALTER ANY SEQUENCE system privilege.

### Keywords and Parameters

The keywords and parameters in this statement serve the same purposes described in ["CREATE SEQUENCE"](#) on page 4-371. In addition:

- To restart the sequence at a different number, you must drop and re-create it (see "[DROP SEQUENCE](#)" on page 7-497).
- If you change the `INCREMENT BY` value *before the first invocation of `NEXTVAL`*, some sequence numbers will be skipped. Therefore, if you want to retain the original `START WITH` value, you must drop the sequence and re-create it with the original `START WITH` value and the new `INCREMENT BY` value.
- Oracle performs some validations. For example, a new `MAXVALUE` cannot be imposed that is less than the current sequence number.

### Examples

This statement sets a new maximum value for the `ESEQ` sequence:

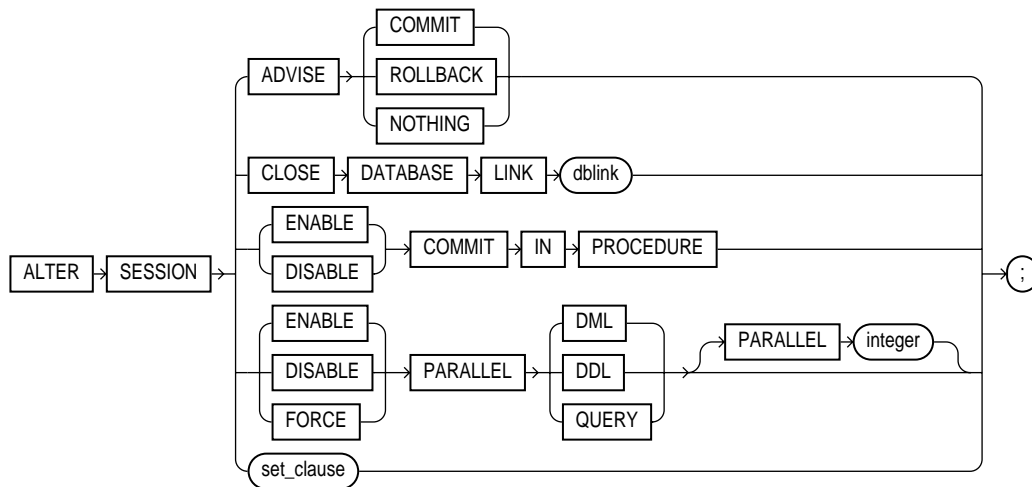
```
ALTER SEQUENCE eseq
  MAXVALUE 1500;
```

This statement turns on `CYCLE` and `CACHE` for the `ESEQ` sequence:


```
ALTER SEQUENCE eseq
  CYCLE
  CACHE 5;
```

## ALTER SESSION

### Syntax



#### set\_clause::=



### Purpose

To specify or modify any of the conditions or parameters that affect your connection to the database. The statement stays in effect until you disconnect from the database.

### Prerequisites

To enable and disable the SQL trace facility, you must have ALTER SESSION system privilege.

You do not need any privileges to perform the other operations of this statement unless otherwise indicated.

---

## Keywords and Parameters

ADVISE	<p>sends advice to a remote database to force a distributed transaction. The advice appears in the <code>ADVISE</code> column of the <code>DBA_2PC_PENDING</code> view on the remote database (the value 'C' for COMMIT, 'R' for ROLLBACK, and ' ' for NOTHING). If the transaction becomes in doubt, the administrator of that database can use this advice to decide whether to commit or roll back the transaction.</p> <p>You can send different advice to different remote databases by issuing multiple <code>ALTER SESSION</code> statements with the <code>ADVISE</code> clause in a single transaction. Each such statement sends advice to the databases referenced in the following statements in the transaction until another such statement is issued. For more information on distributed transactions and how to decide whether to commit or roll back in-doubt distributed transactions, see <i>Oracle8i Distributed Database Systems</i>.</p>
CLOSE DATABASE LINK	<p>closes the database link <i>dblink</i>. When you issue a statement that uses a database link, Oracle creates a session for you on the remote database using that link. The connection remains open until you end your local session or until the number of database links for your session exceeds the value of the initialization parameter <code>OPEN_LINKS</code>. If you want to reduce the network overhead associated with keeping the link open, use this clause to close the link explicitly if you do not plan to use it again in your session. <b>You must first close all cursors that use the link and then end your current transaction if it uses the link.</b></p>
ENABLE   DISABLE COMMIT IN PROCEDURE	<p>Procedures and stored functions written in PL/SQL can issue <code>COMMIT</code> and <code>ROLLBACK</code> statements. If your application would be disrupted by a <code>COMMIT</code> or <code>ROLLBACK</code> statement not issued directly by the application itself, use the <code>DISABLE</code> form of this clause to prevent procedures and stored functions called during your session from issuing these statements.</p> <p>You can subsequently allow procedures and stored functions to issue <code>COMMIT</code> and <code>ROLLBACK</code> statements in your session by issuing the <code>ENABLE</code> form of this clause.</p> <p>Some applications (such as SQL*Forms) automatically prohibit <code>COMMIT</code> and <code>ROLLBACK</code> statements in procedures and stored functions. Refer to your application documentation.</p> <hr/> <p><b>Note:</b> This statement does not apply to database triggers. Triggers can never issue <code>COMMIT</code> or <code>ROLLBACK</code> statements.</p>
PARALLEL DML   DDL   QUERY	<p>specifies whether all subsequent DML, DDL, or query statements in the session will be considered for parallel execution. This clause enables you to override the degree of parallelism of tables during the current session without changing the tables themselves. <b>You can execute this clause for DML only between committed transactions.</b> Uncommitted transactions must either be committed or rolled back prior to executing this clause for DML.</p>
ENABLE	<p>executes subsequent statements in the session in parallel. This is the default for DDL and query statements.</p>



---

	<ul style="list-style-type: none"> <li>■ <b>DML:</b> executes the session's DML statements in parallel mode if a parallel hint or a parallel clause is specified.</li> <li>■ <b>DDL:</b> executes the session's DDL statements in parallel mode if a parallel clause is specified.</li> <li>■ <b>QUERY:</b> executes the session's queries in parallel mode if a parallel hint or a parallel clause is specified</li> </ul> <p><b>Restriction:</b> You cannot specify the optional <code>PARALLEL integer</code> with <code>ENABLE</code>.</p>
DISABLE	<p>specifies that subsequent statements will be executed serially. This is the default for DML statements.</p> <ul style="list-style-type: none"> <li>■ <b>DML:</b> executes the session's DML statements serially.</li> <li>■ <b>DDL:</b> executes the session's DDL statements serially.</li> <li>■ <b>QUERY:</b> executes the session's queries serially.</li> </ul> <p><b>Restriction:</b> You cannot specify the optional <code>PARALLEL integer</code> with <code>DISABLE</code>.</p>
FORCE	<p>forces parallel execution of subsequent statements in the session. If no parallel clause or hint is specified, then a default degree of parallelism is used. This clause overrides any <i>parallel_clause</i> specified in subsequent statements in the session, but is overridden by a parallel hint.</p> <ul style="list-style-type: none"> <li>■ <b>DML:</b> provided no parallel DML restrictions are violated, executes subsequent DML statements in the session with the default degree of parallelism, unless a specific degree is specified in this clause.</li> <li>■ <b>DDL:</b> executes subsequent DDL statements in the session with the default degree of parallelism, unless a specific degree is specified in this clause. Resulting database objects will have associated with them the prevailing degree of parallelism.</li> </ul> <p>Using <code>FORCE DDL</code> automatically causes all tables created in this session to be created with a default level of parallelism. The effect is the same as if you had specified the <i>parallel_clause</i> (with default degree) with the <code>CREATE TABLE</code> statement.</p> <ul style="list-style-type: none"> <li>■ <b>QUERY:</b> executes subsequent queries with the default degree of parallelism, unless a specific degree is specified in this clause.</li> <li>■ <b>PARALLEL integer:</b> explicitly specifies a degree of parallelism <ul style="list-style-type: none"> <li>- For force DDL, the degree overrides any parallel clause in subsequent DDL statements.</li> <li>- For force DML and QUERY, the degree overrides the degree currently stored for the table in the data dictionary.</li> <li>- A degree specified in a statement through a hint will override the degree being forced.</li> </ul> </li> </ul>

---

The following types of DML operations are not parallelized regardless of this clause:

- Operations on clustered tables
- Operations with embedded functions that either write or read database or package states
- Operations on tables with triggers that could fire
- Operations on tables or schema objects containing object types, or LONG or LOB datatypes.

For a detailed description of parallel DML features and hints, see *Oracle8i Designing and Tuning for Performance*.

*set\_clause* sets the session parameters that follow. You can set values for multiple parameters in the same *set\_clause*.

---

**CAUTION: Unless otherwise indicated, the parameters described here are initialization parameters, and the descriptions indicate only the general nature of the parameters. Before changing the values of initialization parameters, please refer to their full description in *Oracle8i Reference* or *Oracle8i National Language Support Guide*.**

---

CONSTRAINT{S} = { IMMEDIATE | DEFERRED | DEFAULT }

determines when conditions specified by a deferrable constraint are enforced. CONSTRAINT{S} is a session parameter only, not an initialization parameter.

- IMMEDIATE indicates that the conditions specified by the deferrable constraint are checked immediately after each DML statement. This setting is equivalent to issuing the SET CONSTRAINTS ALL IMMEDIATE statement at the beginning of each transaction in your session. See the IMMEDIATE parameter of "SET CONSTRAINT(S)" on page 7-598.
- DEFERRED indicates that the conditions specified by the deferrable constraint are checked when the transaction is committed. This setting is equivalent to issuing the SET CONSTRAINTS ALL DEFERRED statement at the beginning of each transaction in your session. See the DEFERRED parameter of "SET CONSTRAINT(S)" on page 7-598.
- DEFAULT restores all constraints at the beginning of each transaction to their initial state of DEFERRED or IMMEDIATE.

CREATE\_STORED\_OUTLINES = { TRUE | FALSE | 'category\_name' }

determines whether Oracle should automatically create and store an outline for each query submitted during the session. CREATE\_STORED\_OUTLINES is not an initialization parameter.

- TRUE enables automatic outline creation for subsequent queries in the same session. These outlines receive a unique system-generated name and are stored in the DEFAULT category. If a particular query already has an outline defined for it in the DEFAULT category, that outline will remain and a new outline will not be created.
- FALSE disables automatic outline creation during the session. This is the default.
- *category\_name* has the same behavior as TRUE except that any outline created during the session is stored in the *category\_name* category.

---

`CURRENT_SCHEMA = schema`

changes the current schema of the session to the specified schema. Subsequent unqualified references to schema objects during the session will resolve to objects in the specified schema. The setting persists for the duration of the session or until you issue another `ALTER SESSION SET CURRENT_SCHEMA` statement. `CURRENT_SCHEMA` is a session parameter only, not an initialization parameter.

This setting offers a convenient way to perform operations on objects in a schema other than that of the current user without having to qualify the objects with the schema name. This setting changes the current schema, but it does not change the session user or the current user, nor does it give you any additional system or object privileges for the session. For more information on this parameter, see *Oracle8i Application Developer's Guide - Fundamentals*.

`CURSOR_SHARING = {FORCE | EXACT}`

determines what kind of SQL statements can share the same cursors.

- `EXACT` causes only identical SQL statements to share a cursor.
- `FORCE` forces statements that may differ in some literals, but are otherwise identical, to share a cursor, unless the literals affect the meaning of the statement.

**See Also:** *Oracle8i Designing and Tuning for Performance* for information on setting this parameter in these and other environments.

`DB_BLOCK_CHECKING = {TRUE | FALSE}`

controls whether data block checking is done. The default is `FALSE`.

`DB_FILE_MULTIBLOCK_READ_COUNT = integer`

specifies with *integer* the maximum number of blocks read in one I/O operation during a sequential scan. The default is 8.

`FAST_START_IO_TARGET = integer`

specifies the target number of I/Os (reads and writes) to and from buffer cache that Oracle should perform upon crash or instance recovery. Oracle continuously calculates the actual number of I/Os that would be needed for recovery and compares that number against the target. If the actual number is greater than the target, Oracle attempts to write additional dirty buffers to advance the checkpoint, while minimizing the affect on performance.

For information on how to tune this parameter, see *Oracle8i Designing and Tuning for Performance*.

`FLAGGER = { ENTRY | INTERMEDIATE | FULL | OFF }`

specifies FIPS flagging, which causes an error message to be generated when a SQL statement issued is an extension of ANSI SQL92. `FLAGGER` is a session parameter only, not an initialization parameter.

In Oracle, there is currently no difference between Entry, Intermediate, or Full level flagging. Once flagging is set in a session, a subsequent `ALTER SESSION SET FLAGGER` statement will work, but generates the message, ORA-00097. This allows FIPS flagging to be altered without disconnecting the session. `OFF` turns off flagging.

`GLOBAL_NAMES = { TRUE | FALSE }`

When you start an instance, Oracle determines whether to enforce global name resolution for remote objects accessed in SQL statements based on the value of the initialization parameter `GLOBAL_NAMES`. This parameter enables or disables global name resolution for the duration of the session. `TRUE` enables the enforcement of global names. `FALSE` disables the enforcement of global names. You can also enable or disable global name resolution for your instance with the `GLOBAL_NAMES` parameter of the `ALTER SYSTEM` statement.

Oracle recommends that you enable global name resolution if you use or plan to use distributed processing. For more information on global name resolution and how Oracle enforces it, see "[Referring to Objects in Remote Databases](#)" on page 2-79 and *Oracle8i Distributed Database Systems*.

`HASH_AREA_SIZE = integer`

specifies in bytes the amount of memory to use for hash join operations. The default is twice the value of the `SORT_AREA_SIZE` initialization parameter.

`HASH_JOIN_ENABLED = {TRUE | FALSE}`

enables or disables the use of the hash join operation in queries. The default is `TRUE`, which enables hash joins.

`HASH_MULTIBLOCK_IO_COUNT = integer`

specifies the number of data blocks to read and write during a hash join operation. The value multiplied by the `DB_BLOCK_SIZE` initialization parameter should not exceed 64 K. The default value for this parameter is 1. If the multi-threaded server is used, the value is always 1, and any value specified here is ignored.

`INSTANCE = integer`

in a parallel server, accesses database files as if the session were connected to the instance specified by *integer*. `INSTANCE` is a session parameter only, not an initialization parameter. For optimum performance, each instance of a parallel server uses its own private rollback segments, freelist groups, and so on. In a parallel server, you normally connect to a particular instance and access data that is partitioned primarily for your use. If you must connect to another instance, the data partitioning can be lost. Setting this parameter lets you access an instance as if you were connected to your own instance.

---

```
ISOLATION_LEVEL = { SERIALIZABLE | READ COMMITTED }
```

specifies how transactions containing database modifications are handled. `ISOLATION_LEVEL` is a session parameter only, not an initialization parameter.

- `SERIALIZABLE` indicates that transactions in the session use the serializable transaction isolation mode as specified in SQL92. That is, if a serializable transaction attempts to execute a DML statement that updates rows currently being updated by another uncommitted transaction at the start of the serializable transaction, then the DML statement fails. A serializable transaction can see its own updates.
- `READ COMMITTED` indicates that transactions in the session will use the default Oracle transaction behavior. Thus, if the transaction contains DML that requires row locks held by another transaction, then the DML statement will wait until the row locks are released.

```
LOG_ARCHIVE_DEST_n = {null_string
                     | {LOCATION=local_pathname | SERVICE=tnsnames_service}
                     [MANDATORY | OPTIONAL] [REOPEN[=integer]]}
```

specifies up to five session-specific valid operating system pathnames or Oracle service names (plus other related options) as destinations for archive redo log file groups ( $n =$  integers 1 through 5). For a description of the options, refer to *Oracle8i Reference*.

**Restrictions:** If you set a value for this parameter,

- You cannot have definitions for the parameters `LOG_ARCHIVE_DEST` and `LOG_ARCHIVE_DUPLEX_DEST` in your initialization parameter file, nor can you set values for those parameters with the `ALTER SYSTEM` statement.
- You cannot start archiving to a specific location using the `ALTER SYSTEM ARCHIVE LOG TO location` statement.

```
LOG_ARCHIVE_DEST_STATE_n = {ENABLE | DEFER}
```

specifies the session-specific state associated with the corresponding `LOG_ARCHIVE_DEST_n` parameter.

- `ENABLE` specifies that any associated valid destination can be used for archiving. This is the default.
- `DEFER` specifies that Oracle will not consider for archiving any destination associated with the corresponding `LOG_ARCHIVE_DEST_n` parameter.

```
LOG_ARCHIVE_MIN_SUCCEED_DEST = integer
```

specifies the session-specific minimum number of destinations that must succeed in order for the online log file to be available for reuse.

```
MAX_DUMP_FILE_SIZE = { size | UNLIMITED }
```

specifies the upper limit of trace dump file size. Specify the maximum *size* as either a nonnegative integer that represents the number of blocks, or as `UNLIMITED`. If you specify `UNLIMITED`, no upper limit is imposed.

**NLS parameters:** When you start an instance, Oracle establishes support based on the values of initialization parameters that begin with "NLS". You can query the dynamic performance table `V$NLS_PARAMETERS` to see the current NLS attributes for your session. For more information about NLS parameters, see *Oracle8i National Language Support Guide*.

```
NLS_CALENDAR = 'text'
```

explicitly specifies a new calendar type.

```
NLS_COMP = 'text'
```

specifies that linguistic comparison is to be used according to the `NLS_SORT` parameter. This parameter obviates the need to specify `NLS_SORT` in SQL statements.

```
NLS_CURRENCY = 'text'
```

explicitly specifies a new value for the L number format element (the local currency symbol). The symbol cannot exceed 10 characters.

```
NLS_DATE_FORMAT = 'fmt'
```

explicitly specifies a new default date format. The *fmt* value must be a date format model as specified in the section "[Date Format Elements](#)" on page 2-45.

```
NLS_DATE_LANGUAGE = language
```

explicitly changes the language for names and abbreviations of days and months, and for spelled-out values of other date format elements.

```
NLS_DUAL_CURRENCY = 'text'
```

explicitly specifies a new "Euro" (or other) dual currency symbol. The value of *text* is returned by the number format element U (see "[Number Format Elements](#)" on page 2-41); *text* cannot exceed 10 characters.

```
NLS_ISO_CURRENCY = territory
```

explicitly specifies the territory whose ISO currency symbol should be used. That territory's currency symbol then becomes the value of the C number format element.

```
NLS_LANGUAGE = language
```

changes the language in which Oracle returns errors and other messages. This parameter also implicitly specifies new values for these items:

- Language for day and month names and abbreviations and spelled values of other elements
- Linguistic sort sequences or binary sorts
- B.C. and A.D. indicators
- A.M. and P.M. meridian indicators

`NLS_NUMERIC_CHARACTERS = 'text'`

explicitly specifies a new decimal character and group separator. The *'text'* value must have this form:

`'dg'`

where: *d* is the new decimal character, and *g* is the new group separator.

The decimal character and the group separator must be two different single-byte characters, and cannot be a numeric value or any of the following characters: plus sign ("+"), minus sign or hyphen ("-"), less-than sign("<"), or greater-than sign(">").

If the decimal character is not a period (.), you must use single quotation marks to enclose all number values that appear in expressions in your SQL statements. When not using a period for the decimal point, use the `TO_NUMBER` function to ensure that a valid number is retrieved.

`NLS_SORT = { sort | BINARY }`

changes the sequence into which Oracle sorts character values. *sort* specifies the name of a linguistic sort sequence. `BINARY` specifies a binary sort. The default is `BINARY`.

`NLS_TERRITORY = territory`

implicitly specifies new values for these items:

- Default date format
- Decimal character and group separators
- Local currency symbol
- ISO currency symbol
- First day of the week for D date format element

`OBJECT_CACHE_MAX_SIZE_PERCENT = integer`

specifies the percentage of the optimal cache size that the session object cache can grow beyond the optimal size. The default is 10.

`OBJECT_CACHE_OPTIMAL_SIZE = integer`

specifies (in kilobytes) the size to which the session object cache is reduced when it exceeds maximum size. The default is 100.

`OPTIMIZER_INDEX_CACHING = integer`

lets you tune the optimizer to favor nested loops joins and IN-list iterators. The value of *integer* indicates the percentage of the index blocks assumed to be in the cache.

`OPTIMIZER_INDEX_COST_ADJ = integer`

let you tune optimizer behavior for access path selection to make the optimizer more likely to select an index access path than a full table scan. The value of *integer* is a percentage indicating the importance the optimizer attaches to the index path compared with "normal". The default is 100 (indicating 100%), which makes the optimizer cost index access paths at the regular cost.

`OPTIMIZER_MAX_PERMUTATIONS = integer`

lets you limit the amount of work the optimizer expends on optimizing queries with large joins. The value of *integer* is the number of permutations of the tables the optimizer will consider with large joins.

`OPTIMIZER_MODE = { ALL_ROWS | FIRST_ROWS | RULE | CHOOSE }`

specifies the approach and mode of the optimizer for your session. For information on how to choose a goal for the cost-based approach based on the characteristics of your application, see *Oracle8i Concepts* and *Oracle8i Designing and Tuning for Performance*.

- `ALL_ROWS` specifies the cost-based approach and optimizes for best throughput.
- `FIRST_ROWS` specifies the cost-based approach and optimizes for best response time.
- `RULE` specifies the rule-based approach. (The rule-based optimizer does not use function-based indexes.)
- `CHOOSE` causes the optimizer to choose an optimization approach based on the presence of statistics in the data dictionary.

`OPTIMIZER_PERCENT_PARALLEL = integer`

specifies the amount of parallelism the optimizer uses in its cost functions. The default is 0 (no parallelism).

`PARALLEL_BROADCAST_ENABLED = { TRUE | FALSE }`

lets you enhance performance during hash and merge joins.

`PARALLEL_INSTANCE_GROUP = ' text '`

identifies the parallel instance group to be used for spawning parallel query slaves. The default is all active instances. **Set this parameter only if you are running Oracle Parallel Server in parallel mode.**

`PARALLEL_MIN_PERCENT = integer`

specifies the minimum percent of threads required for parallel query. The default is 0 (no parallelism).



---

PARTITION\_VIEW\_ENABLED = { TRUE | FALSE }

When set to TRUE, this parameter causes the optimizer to skip unnecessary table accesses in a partition view.

**Note:** For important information on partition views, see "[Partition Views](#)" on page 7-457.

---

PLSQL\_V2\_COMPATIBILITY = { TRUE | FALSE }

if TRUE, modifies the compile-time behavior of PL/SQL programs to allow language constructs that are illegal in Oracle8 and Oracle8i (PL/SQL V3), but were legal in Oracle7 (PL/SQL V2). FALSE disallows illegal Oracle7 PL/SQL V2 constructs. This is the default.

See the *PL/SQL User's Guide and Reference* and *Oracle8i Reference* for more information about this session parameter.

QUERY\_REWRITE\_ENABLED = { TRUE | FALSE }

enables or disables query rewrite on all materialized views that have not been explicitly disabled. Query rewrite is disabled by default. It is also disabled by rule-based optimization (that is, if the OPTIMIZER\_MODE parameter is set to RULE). For more information on query rewrite, see *Oracle8i Data Warehousing Guide*.

This parameter has the following additional effect on the use of function-based indexes:

- If this parameter is set to TRUE, Oracle will use function-based indexes to derive values of SQL expressions. If in addition the QUERY\_REWRITE\_INTEGRITY parameter is set to any value other than ENFORCED, Oracle will derive such values even if the index is based on a user-defined (rather than SQL) function.
- If this parameter is set to FALSE, Oracle will not use function-based indexes to derive values of SQL expressions, but it will use such indexes to obtain values of real columns in the index.

Enabling or disabling query rewrite does not affect descending indexes.

A setting of TRUE has no effect on materialized views that cannot be created with the ENABLE QUERY REWRITE clause, such as materialized views created totally or in part from a view.

QUERY\_REWRITE\_INTEGRITY = { ENFORCED | TRUSTED | STALE\_TOLERATED }

sets the minimum consistency level for query rewrite. The following values are permitted:

- ENFORCED is the safest level. It relies only on system-enforced relationships so that data integrity and correctness can be guaranteed. This level ensures that query rewrite will not use any function-based index or any materialized view that includes a call to a user-defined function.

In addition, this level ensures that query rewrite will not use any dimensional information or any constraints enabled with the RELY keyword.

- TRUSTED specifies that materialized views created with the ON PREBUILT TABLE clause are supported, and trusted but unenforced join relationships are accepted. Query rewrite uses join information from dimensions and enables unenforced constraints with the RELY keyword.
- STALE\_TOLERATED specifies that any stale, usable materialized view may be used.

This parameter does not affect descending indexes.

For more information on query rewrite integrity level, see *Oracle8i Data Warehousing Guide*. For information on dimensions, see "CREATE DIMENSION" on page 7-277. For information on constraints enabled with the RELY keyword, see "constraint\_clause" on page 7-233.

```
REMOTE_DEPENDENCIES_MODE = { TIMESTAMP | SIGNATURE }
```

specifies how dependencies of remote stored procedures are handled by the session. For more information, refer to *Oracle8i Application Developer's Guide - Fundamentals*.

```
SESSION_CACHED_CURSORS = integer
```

specifies the number of frequently used cursors that can be retained in the cache. The cursors can be open or closed, which is particularly useful for Oracle tools that close all session cursors associated with a form when switching to another form. In such cases, frequently used cursors do not have to be reparsed. A least recently used algorithm ages out entries in the cache to make room for new entries when needed. For more information on session cursor caching, see *Oracle8i Designing and Tuning for Performance*.

```
SKIP_UNUSABLE_INDEXES = { TRUE | FALSE }
```

controls the use and reporting of tables with unusable indexes or index partitions. SKIP\_UNUSABLE\_INDEXES is a session parameter only, not an initialization parameter.

- TRUE disables error reporting of indexes and index partitions marked UNUSABLE. Allows all operations (inserts, deletes, updates, and selects) to tables with unusable indexes or index partitions.

---

**Note:** Statements that would normally use the unusable indexes or index partitions may be compiled with suboptimal optimizer plans, occasionally resulting in major degradation in response time and resource utilization.

---

- FALSE enables error reporting of indexes marked UNUSABLE. Does not allow inserts, deletes, and updates to tables with unusable indexes or index partitions. This is the default.

---

```
SORT_AREA_RETAINED_SIZE = integer
```

specifies (in bytes) the maximum amount of memory that each sort operation will retain after the first fetch is done, until the cursor ends. If you do not explicitly set this parameter in the initialization parameter file or dynamically, Oracle uses the value of the `SORT_AREA_SIZE` parameter.

```
SORT_AREA_SIZE = integer
```

specifies (in bytes) the maximum amount of memory to use for each sort operation. The default is OS-dependent.

```
SORT_MULTIBLOCK_READ_COUNT = integer
```

specifies the number of database blocks to read each time a sort performs a read from temporary segments. The default is 2.

```
SQL_TRACE = { TRUE | FALSE }
```

The SQL trace facility generates performance statistics for the processing of SQL statements. When you begin a session, Oracle enables or disables the SQL trace facility based on the value of this parameter. You can subsequently enable or disable the SQL trace facility for your own session with the `SQL_TRACE` parameter of the `ALTER SESSION` statement. `TRUE` enables the SQL trace facility. `FALSE` disables it.

`SQL_TRACE` is an initialization parameter. However, when you change its value with an `ALTER SESSION` statement, the results are not reflected in the `V$PARAMETER` view. Therefore, in this context it is considered a session parameter only.

**See Also:** *Oracle8i Designing and Tuning for Performance* for more information on the SQL trace facility, including how to format and interpret its output.

```
STAR_TRANSFORMATION_ENABLED = { TRUE | FALSE }
```

determines whether a cost-based query transformation will be applied to star queries. The default is `FALSE`.

```
TIMED_STATISTICS = { TRUE | FALSE }
```

specifies whether the server requests the time from the operating system when generating time-related statistics. The default is `FALSE`.

```
USE_STORED_OUTLINES = { TRUE | FALSE | 'category_name' }
```

determines whether the optimizer will use stored outlines to generate execution plans. `USE_STORED_OUTLINES` is not an initialization parameter.

- `TRUE` causes the optimizer to use outlines stored in the `DEFAULT` category when compiling requests.
  - `FALSE` specifies that the optimizer should not use stored outlines. This is the default.
  - `category_name` causes the optimizer to use outlines stored in the `category_name` category when compiling requests.
-

## Examples

**PARALLEL Example** Issue the following statement to enable parallel DML mode for the current session:

```
ALTER SESSION ENABLE PARALLEL DML;
```

**ADVISE Example** The following transaction inserts an employee record into the EMP table on the database identified by the database link SITE1 and deletes an employee record from the EMP table on the database identified by SITE2:

```
ALTER SESSION
  ADVISE COMMIT;

INSERT INTO emp@site1
  VALUES (8002, 'FERNANDEZ', 'ANALYST', 7566,
    TO_DATE('04-OCT-1992', 'DD-MON-YYYY'), 3000, NULL, 20);

ALTER SESSION
  ADVISE ROLLBACK;

DELETE FROM emp@site2
  WHERE empno = 8002;

COMMIT;
```

This transaction has two ALTER SESSION statements with the ADVISE clause. If the transaction becomes in doubt, SITE1 is sent the advice 'COMMIT' by virtue of the first ALTER SESSION statement and SITE2 is sent the advice 'ROLLBACK' by virtue of the second.

**CLOSE DATABASE LINK Example** This statement updates the employee table on the SALES database using a database link, commits the transaction, and explicitly closes the database link:

```
UPDATE emp@sales
  SET sal = sal + 200
  WHERE empno = 9001;

COMMIT;

ALTER SESSION
  CLOSE DATABASE LINK sales;
```

**Date Format Example** The following statement dynamically changes the default date format for your session to 'YYYY MM DD-HH24:MI:SS':

```
ALTER SESSION
  SET NLS_DATE_FORMAT = 'YYYY MM DD HH24:MI:SS';
```

Oracle uses the new default date format:

```
SELECT TO_CHAR(SYSDATE) Today
  FROM DUAL;
```

```
TODAY
-----
1997 08 12 14:25:56
```

**Date Language Example** The following statement changes the language for date format elements to French:

```
ALTER SESSION
  SET NLS_DATE_LANGUAGE = French;

SELECT TO_CHAR(SYSDATE, 'Day DD Month YYYY') Today
  FROM DUAL;
```

```
TODAY
-----
Mardi      28 Février    1997
```

**ISO Currency Example** The following statement dynamically changes the ISO currency symbol to the ISO currency symbol for the territory America:

```
ALTER SESSION
  SET NLS_ISO_CURRENCY = America;

SELECT TO_CHAR( SUM(sal), 'C999G999D99') Total
  FROM emp;
```

```
TOTAL
-----
USD29,025.00
```

**Decimal Character and Group Separator Example** The following statement dynamically changes the decimal character to comma (,) and the group separator to period (.):

```
ALTER SESSION SET NLS_NUMERIC_CHARACTERS = ',.' ;
```

Oracle returns these new characters when you use their number format elements:

```
SELECT TO_CHAR( SUM(sal), 'L999G999D99') Total FROM emp ;
```

```
TOTAL
-----
FF29.025,00
```

**NLS Currency Example** The following statement dynamically changes the local currency symbol to 'DM':

```
ALTER SESSION
  SET NLS_CURRENCY = 'DM';

SELECT TO_CHAR( SUM(sal), 'L999G999D99') Total
       FROM emp;
```

```
TOTAL
-----
DM29.025,00
```

**NLS Language Example** The following statement dynamically changes to French the language in which error messages are displayed:

```
ALTER SESSION
  SET NLS_LANGUAGE = FRENCH;

SELECT * FROM DMP;

ORA-00942: Table ou vue inexistante
```

**Linguistic Sort Example** The following statement dynamically changes the linguistic sort sequence to Spanish:

```
ALTER SESSION
  SET NLS_SORT = XSpanish;
```

Oracle sorts character values based on their position in the Spanish linguistic sort sequence.

**SQL Trace Example** To enable the SQL trace facility for your session, issue the following statement:

```
ALTER SESSION
  SET SQL_TRACE = TRUE;
```

**Query Rewrite Example** This statement enables query rewrite in the current session for all materialized views that have not been explicitly disabled:

```
ALTER SESSION SET QUERY_REWRITE_ENABLED = TRUE;
```

---

## ALTER SNAPSHOT

In Oracle8i, "snapshots" are synonymous with "materialized views." Please see "[ALTER MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-47.

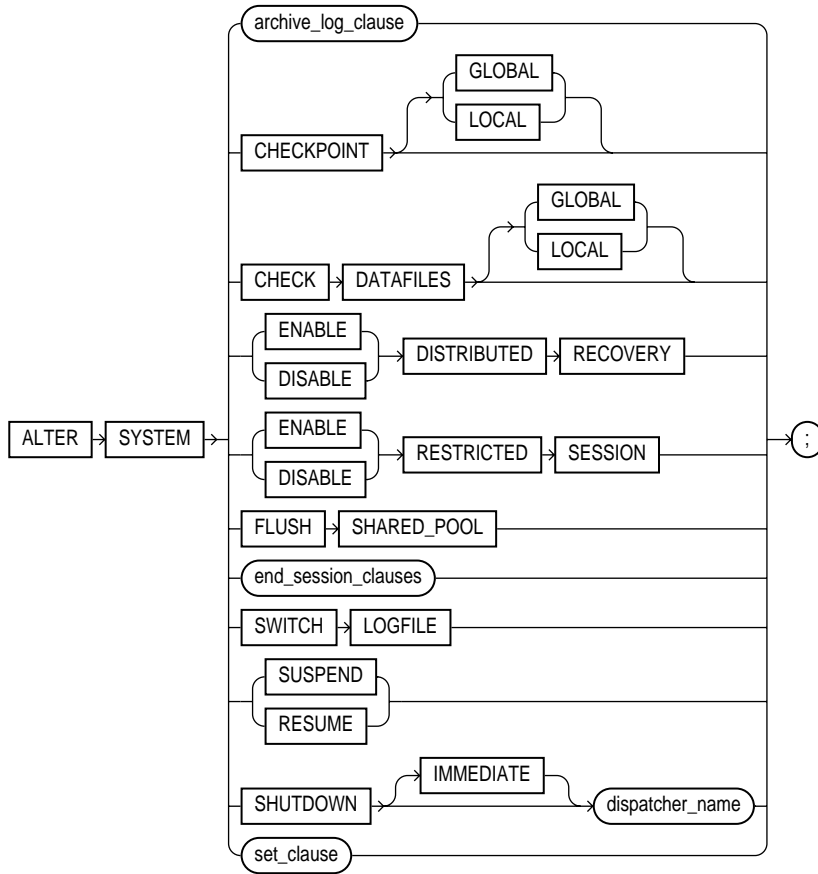


## ALTER SNAPSHOT LOG


In Oracle8i, "snapshots" are synonymous with "materialized views." Please see "[ALTER MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-58.

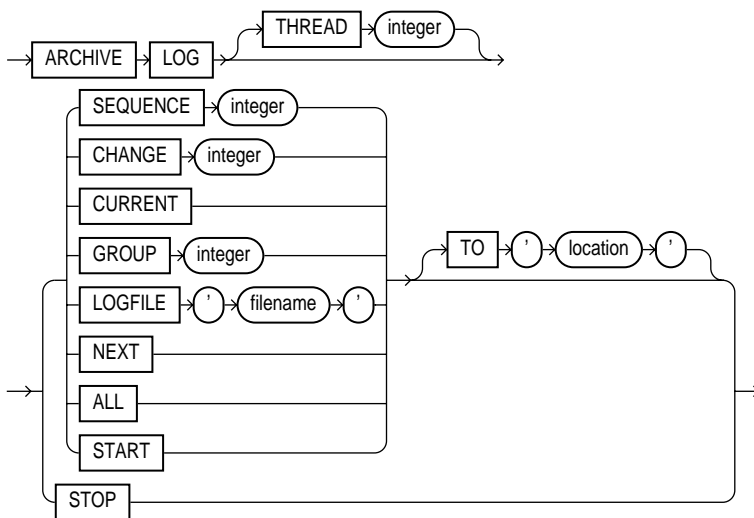

# ALTER SYSTEM

## Syntax



### end\_session\_clauses::=



**archive\_log\_clause::=****set\_clause::=****Purpose**

To dynamically alter your Oracle instance. The settings stay in effect as long as the database is mounted.

**Prerequisites**

You must have ALTER SYSTEM system privilege.

To specify the *archive\_log\_clause*, you must have the OSDBA or OSOPER role enabled.

**Keywords and Parameters**


---

<i>archive_log_clause</i>	manually archives redo log files or enables or disables automatic archiving. To use this clause, your instance must have the database mounted. The database can be either open or closed unless otherwise noted.
---------------------------	--

**Notes:**

- You can also manually archive redo log file groups with the `ARCHIVE LOG SQL*Plus` statement. For information on this statement, see the *SQL\*Plus User's Guide and Reference*.
- You can also have Oracle archive redo log files groups automatically. For information on automatic archiving, see *Oracle8i Administrator's Guide*. You can always manually archive redo log file groups regardless of whether automatic archiving is enabled.

---

THREAD	specifies the thread containing the redo log file group to be archived. <b>Set this parameter only if you are using Oracle with the Parallel Server option in parallel mode.</b>
SEQUENCE	manually archives the online redo log file group identified by the log sequence number <i>integer</i> in the specified thread. If you omit the <code>THREAD</code> parameter, Oracle archives the specified group from the thread assigned to your instance.
CHANGE	manually archives the online redo log file group containing the redo log entry with the system change number (SCN) specified by <i>integer</i> in the specified thread. If the SCN is in the current redo log file group, Oracle performs a log switch. If you omit the <code>THREAD</code> parameter, Oracle archives the groups containing this SCN from all enabled threads. <b>You can use this clause only when your instance has the database open.</b>
CURRENT	manually archives the current redo log file group of the specified thread, forcing a log switch. If you omit the <code>THREAD</code> parameter, Oracle archives all redo log file groups from all enabled threads, including logs previous to current logs. <b>You can use this clause only when your instance has the database open.</b>

---

**Note:** If you specify a redo log file group for archiving with the `CHANGE` or `CURRENT` clause, and earlier redo log file groups are not yet archived, Oracle archives all unarchived groups up to and including the specified group.

---

GROUP	manually archives the online redo log file group with the <code>GROUP</code> value specified by <i>integer</i> . You can determine the <code>GROUP</code> value for a redo log file group by examining the data dictionary view <code>DBA_LOG_FILES</code> . If you specify both the <code>THREAD</code> and <code>GROUP</code> parameters, the specified redo log file group must be in the specified thread.
LOGFILE	manually archives the online redo log file group containing the redo log file member identified by ' <i>filename</i> '. If you specify both the <code>THREAD</code> and <code>LOGFILE</code> parameters, the specified redo log file group must be in the specified thread.

**Restriction:** You must archive redo log file groups in the order in which they are filled. If you specify a redo log file group for archiving with the `LOGFILE` parameter, and earlier redo log file groups are not yet archived, Oracle returns an error.

---

NEXT	manually archives the next online redo log file group from the specified thread that is full but has not yet been archived. If you omit the <code>THREAD</code> parameter, Oracle archives the earliest unarchived redo log file group from any enabled thread.
ALL	manually archives all online redo log file groups from the specified thread that are full but have not been archived. If you omit the <code>THREAD</code> parameter, Oracle archives all full unarchived redo log file groups from all enabled threads.
START	enables automatic archiving of redo log file groups. <b>Restriction:</b> You can enable automatic archiving only for the thread assigned to your instance.
TO ' <i>location</i> '	specifies the primary location to which the redo log file groups are archived. The value of this parameter must be a fully specified file location following the conventions of your operating system. If you omit this parameter, Oracle archives the redo log file group to the location specified by the initialization parameters <code>LOG_ARCHIVE_DEST</code> or <code>LOG_ARCHIVE_DEST_n</code> . <hr/> <b>Note:</b> You can enhance recovery reliability by setting the related archive parameters <code>LOG_ARCHIVE_DEST_DUPLEX</code> and <code>LOG_ARCHIVE_MIN_SUCCEED_DEST</code> . <hr/>
STOP	disables automatic archiving of redo log file groups. You can disable automatic archiving only for the thread assigned to your instance.
CHECKPOINT	explicitly forces Oracle to perform a checkpoint, ensuring that all changes made by committed transactions are written to datafiles on disk. You can specify this clause only when your instance has the database open. Oracle does not return control to you until the checkpoint is complete.
GLOBAL	in an Oracle Parallel Server environment, performs a checkpoint for all instances that have opened the database. This is the default.
LOCAL	in an Oracle Parallel Server environment, performs a checkpoint only for the thread of redo log file groups for your instance.
	<b>See Also:</b> <i>Oracle8i Concepts</i> for more information on checkpoints.
CHECK DATAFILES	in a distributed database system, such as an Oracle Parallel Server environment, updates an instance's SGA from the database control file to reflect information on all online datafiles.
GLOBAL	performs this synchronization for all instances that have opened the database. This is the default.
LOCAL	performs this synchronization only for the local instance.
	Your instance should have the database open.
	<b>See Also:</b> <i>Oracle8i Parallel Server Setup and Configuration Guide</i> .

---

DISCONNECT SESSION	<p>disconnects the current session by destroying the dedicated server process (or virtual circuit if the connection was made by way of a multi-threaded server). To use this clause, your instance must have the database open. You must identify the session with both of the following values from the V\$SESSION view:</p> <p><i>integer1</i> is the value of the SID column.</p> <p><i>integer2</i> is the value of the SERIAL# column.</p> <p>If system parameters are appropriately configured, application failover will take effect.</p> <p><b>See Also:</b> <i>Oracle8i Parallel Server Setup and Configuration Guide</i> for more information about application failover.</p>
POST_ TRANSACTION	<p>allows ongoing transactions to complete before the session is disconnected. If the session has no ongoing transactions, this clause has the same effect as <b>KILL SESSION</b>, described below.</p>
IMMEDIATE	<p>disconnects the session and recovers the entire session state immediately, without waiting for ongoing transactions to complete.</p> <ul style="list-style-type: none"> <li>■ If you also specify <b>POST_TRANSACTION</b> and the session has ongoing transactions, the <b>IMMEDIATE</b> keyword is ignored.</li> <li>■ If you do not specify <b>POST_TRANSACTION</b>, or you specify <b>POST_TRANSACTION</b> but the session has no ongoing transactions, this clause has the same effect as <b>KILL SESSION IMMEDIATE</b>, described below.</li> </ul>
DISTRIBUTED RECOVERY	<p>specifies whether or not distributed recovery is enabled. To use this clause, your instance must have the database open.</p> <p><b>ENABLE</b> enables distributed recovery. In a single-process environment, you must use this clause to initiate distributed recovery.</p> <p>You may need to issue the <b>ENABLE DISTRIBUTED RECOVERY</b> statement more than once to recover an in-doubt transaction if the remote node involved in the transaction is not accessible. In-doubt transactions appear in the data dictionary view <b>DBA_2PC_PENDING</b>.</p> <p><b>See Also:</b> <i>Oracle8i Distributed Database Systems</i> for more information about distributed transactions and distributed recovery.</p>
	<p><b>DISABLE</b> disables distributed recovery.</p>
RESTRICTED SESSION	<p>specifies whether logon to Oracle is restricted</p> <p><b>ENABLE</b> allows only users with <b>RESTRICTED SESSION</b> system privilege to log on to Oracle. Existing sessions are not terminated.</p> <p><b>DISABLE</b> reverses the effect of the <b>ENABLE RESTRICTED SESSION</b> clause, allowing all users with <b>CREATE SESSION</b> system privilege to log on to Oracle. This is the default.</p>

---

	<p>You can use this clause regardless of whether your instance has the database dismounted or mounted, open or closed.</p>
FLUSH SHARED_POOL	<p>clears all data from the shared pool in the system global area (SGA). The shared pool stores</p> <ul style="list-style-type: none"> <li>■ Cached data dictionary information and</li> <li>■ Shared SQL and PL/SQL areas for SQL statements, stored procedures, function, packages, and triggers.</li> </ul> <p>This statement does not clear shared SQL and PL/SQL areas for items that are currently being executed. You can use this clause regardless of whether your instance has the database dismounted or mounted, open or closed.</p>
KILL SESSION	<p>marks a session as dead, rolls back ongoing transactions, releases all session locks, and partially recovers session resources. To use this clause, your instance must have the database open, and your session and the session to be killed must be on the same instance. You must identify the session with both of the following values from the V\$SESSION view:</p> <p><i>integer1</i>            is the value of the SID column.</p> <p><i>integer2</i>            is the value of the SERIAL# column.</p> <p>If the session is performing some activity that must be completed, such as waiting for a reply from a remote database or rolling back a transaction, Oracle waits for this activity to complete, marks the session as dead, and then returns control to you. If the waiting lasts a minute, Oracle marks the session to be killed and returns control to you with a message that the session is marked to be killed. The PMON background process then marks the session as dead when the activity is complete.</p> <p>Whether or not the session has an ongoing transaction, Oracle does not recover the entire session state until the session user issues a request to the session and receives a message that the session has been killed.</p> <p>IMMEDIATE            rolls back ongoing transactions, releases all session locks, recovers the entire session state, and returns control to you immediately.</p>
SWITCH LOGFILE	<p>explicitly forces Oracle to begin writing to a new redo log file group, regardless of whether the files in the current redo log file group are full. When you force a log switch, Oracle begins to perform a checkpoint. Oracle returns control to you immediately rather than when the checkpoint is complete. To use this clause, your instance must have the database open.</p>
SUSPEND	<p>suspends all I/O (datafile, control file, and file header) as well as queries, in all instances, enabling you to make copies of the database without having to handle ongoing transactions.</p>

**Restrictions:**

- Do not use this clause unless you have put the database tablespaces in hot backup mode.
- If you start a new instance while the system is suspended, that new instance will not be suspended.

---

RESUME	<p>makes the database available once again for queries and I/O.</p> <p><b>See Also:</b> <i>Oracle8i Backup and Recovery Guide</i> for more information on the <code>SUSPEND</code> clause and <code>RESUME</code> clause.</p>
SHUTDOWN	<p>This clause is relevant only if your system is using Oracle's multi-threaded server architecture. It shuts down a dispatcher identified by <i>dispatcher_name</i>. The <i>dispatcher_name</i> must be a string of the form 'Dxxx', where xxx indicates the number of the dispatcher. (For a listing of dispatcher names, query the <code>NAME</code> column of the <code>V\$DISPATCHER</code> dynamic performance view.)</p> <ul style="list-style-type: none"> <li>■ If you specify <code>IMMEDIATE</code>, the dispatcher stops accepting new connections immediately and Oracle terminates all existing connections through that dispatcher. After all sessions are cleaned up, the dispatcher process literally shuts down.</li> <li>■ If you do not specify <code>IMMEDIATE</code>, the dispatcher stops accepting new connections immediately but waits for all its users to disconnect and for all its database links to terminate. Then it literally shuts down.</li> </ul> <p><b>See Also:</b> <i>Oracle8i Administrator's Guide</i>, <i>Net8 Administrator's Guide</i>, and <i>Oracle8i Designing and Tuning for Performance</i> for more information on dispatchers and multi-threaded server architecture.</p>
<i>set_clause</i>	<p>sets the system parameters that follow. You can set values for multiple parameters in the same <i>set_clause</i>. <b>The <code>DEFERRED</code> keyword sets or modifies the value of the parameter for future sessions that connect to the database.</b></p> <hr/> <p><b>CAUTION:</b> Unless otherwise noted, these parameters are initialization parameters, and the descriptions provided here indicate only the general nature of the parameters. Before changing the values of initialization parameters, please refer to their full description in <i>Oracle8i Reference</i> and <i>Oracle8i National Language Support Guide</i>.</p> <hr/>
<code>AQ_TM_PROCESSES = integer</code>	<p>is an Advanced Queuing parameter that specifies whether a time manager process is created. Accepted values are 1 (creates one time manager process to monitor messages) and 0 (does not create a time manager process).</p>
<code>BACKGROUND_DUMP_DEST = 'text'</code>	<p>specifies the pathname for a directory where debugging trace files for the background processes are written during Oracle operations.</p>
<code>BACKUP_TAPE_IO_SLAVES = {TRUE   FALSE} DEFERRED</code>	<p>specifies whether I/O slaves are used by the Recovery Manager to back up, copy, or restore data to tape.</p>
<code>CONTROL_FILE_RECORD_KEEP_TIME = integer</code>	<p>specifies the minimum of days before a reusable record in the control file can be reused.</p>
<code>CORE_DUMP_DEST = 'text'</code>	



---

specifies the directory where Oracle dumps core files.

`CREATE_STORED_OUTLINES = { TRUE | FALSE | 'category_name' } [NOOVERRIDE]`

determines whether Oracle should automatically create and store an outline for each query submitted on the system. `CREATE_STORED_OUTLINES` is not an initialization parameter.

- `TRUE` enables automatic outline creation for subsequent queries in the system. These outlines receive a unique system-generated name and are stored in the `DEFAULT` category. If a particular query already has an outline defined for it in the `DEFAULT` category, that outline will remain and a new outline will not be created.
- `FALSE` disables automatic outline creation for the system. This is the default.
- `category_name` has the same behavior as `TRUE` except that any outline created in the system is stored in the `category_name` category.
- `NOOVERRIDE` specifies that this system setting will not override the setting for any session in which this parameter was explicitly set. If you do not specify `NOOVERRIDE`, this setting takes effect in all sessions.

`CURSOR_SHARING = { FORCE | EXACT }`

determines what kind of SQL statements can share the same cursors.

- `EXACT` causes only identical SQL statements to share a cursor.
- `FORCE` forces statements that may differ in some literals, but are otherwise identical, to share a cursor, unless the literals affect the meaning of the statement.

**See Also:** for information on setting this parameter in these and other environments, see *Oracle8i Designing and Tuning for Performance*.

`DB_BLOCK_CHECKING = { TRUE | FALSE } DEFERRED`

controls whether data block checking is done. The default is `FALSE`, for compatibility with earlier releases where block checking is disabled as a default.

`DB_BLOCK_CHECKSUM = { TRUE | FALSE }`

specifies whether the database writer background process and the direct loader will calculate a checksum and store it in the cache header of every data lock when writing to disk.

`DB_BLOCK_MAX_DIRTY_TARGET = integer`

limits to *integer* the number of dirty buffers in the cache and reduces the number of buffers that will need to be read during crash or instance recovery. This parameter does **not** relate to media recovery. A value of 0 disables this parameter. The minimum accepted value to enable the parameter is 1000.

---

**Note:** Oracle Corporation recommends that Enterprise Edition users who were using incremental checkpointing in an earlier release now use fast-start checkpointing in Oracle8i. In fast-start checkpointing, the FAST\_START\_IO\_TARGET parameter takes the place of DB\_FILE\_MAX\_DIRTY\_TARGET. See FAST\_START\_IO\_TARGET below. See also *Oracle8i Backup and Recovery Guide* for information on fast-start checkpointing and *Oracle8i Reference* for information on the new parameters.

---

DB\_FILE\_DIRECT\_IO\_COUNT = integer DEFERRED

specifies the number of blocks Oracle should use for I/O during backup, restore, or direct-path read and write operations.

DB\_FILE\_MULTIBLOCK\_READ\_COUNT = integer

specifies the maximum number of blocks read in one I/O operation during a sequential scan.

FAST\_START\_IO\_TARGET = integer

specifies the target number of I/Os (reads and writes) to and from buffer cache that Oracle should perform upon crash or instance recovery. Oracle continuously calculates the actual number of I/Os that would be needed for recovery and compares that number against the target. If the actual number is greater than the target, Oracle attempts to write additional dirty buffers to advance the checkpoint, while minimizing the affect on performance.

For information on how to tune this parameter, see *Oracle8i Designing and Tuning for Performance*.

FAST\_START\_PARALLEL\_ROLLBACK = { FALSE | LOW | HIGH }

specifies the number of processes spawned to perform parallel recovery.

- FALSE specifies no parallel recovery. SMON will serially recover dead transactions.
- LOW specifies that the number of recovery servers may not exceed twice the value of the CPU\_COUNT parameter.
- HIGH specifies that the number of recovery servers may not exceed four times the value of the CPU\_COUNT parameter.

FIXED\_DATE = { 'DD\_MM\_YY' | 'YYYY\_MI\_DD\_HH24\_MI-SS' }

specifies a constant date for SYSDATE instead of the current date.

GC\_DEFER\_TIME = integer

specifies the time (in hundredths of seconds) that Oracle waits before responding to forced-write requests from other instances.

GLOBAL\_NAMES = { TRUE | FALSE }

---

When you start an instance, Oracle determines whether to enforce global name resolution for remote objects accessed in SQL statements based on the value of the initialization parameter `GLOBAL_NAMES`. This system parameter enables or disables global name resolution while your instance is running. `TRUE` enables the enforcement of global names. `FALSE` disables the enforcement of global names. You can also enable or disable global name resolution for your session with the `GLOBAL_NAMES` parameter of the `ALTER SESSION` statement.

Oracle recommends that you enable global name resolution if you use or plan to use distributed processing.

**See Also:** "Referring to Objects in Remote Databases" on page 2-79 and *Oracle8i Distributed Database Systems* for more information on global name resolution and how Oracle enforces it.

`HASH_MULTIBLOCK_IO_COUNT = integer`

specifies the number of data blocks to read and write during a hash join operation. The value multiplied by the `DB_BLOCK_SIZE` initialization parameter should not exceed 64K. The default value for this parameter is 1. If the multi-threaded server is used, the value is always 1, and any value given here is ignored.

`HS_AUTOREGISTER = {TRUE | FALSE}`

enables or disables automatic self-registration of non-Oracle system characteristics in the Oracle server's data dictionary by Heterogeneous Services agents.

**See Also:** *Oracle8i Distributed Database Systems* for more information on accessing non-Oracle systems through Heterogeneous Services.

`JOB_QUEUE_PROCESSES = integer`

specifies the number of job queue processes per instance (SNP $n$ , where  $n$  is 0 to 9 followed by A to Z). Set this parameter to 1 or higher if you wish to have your snapshots updated automatically. One job queue process is usually sufficient unless you have many snapshots that refresh simultaneously.

Oracle also uses job queue processes to process requests created by the `DBMS_JOB` package.

**See Also:** *Oracle8i Replication* for more information on managing table snapshots.

`LICENSE_MAX_SESSIONS = integer`

resets (for the current instance) the value of the initialization parameter `LICENSE_MAX_SESSIONS`, which establishes the concurrent usage licensing limit, or the limit for concurrent sessions. Once this limit is reached, only users with `RESTRICTED SESSION` system privilege can connect. A value of 0 disables the limit.

If you reduce the limit on sessions below the current number of sessions, Oracle does not end existing sessions to enforce the new limit. However, users without `RESTRICTED SESSION` system privilege can begin new sessions only when the number of sessions falls below the new limit.

**Do not disable or raise session limits unless you have appropriately upgraded your Oracle license. For more information, contact your Oracle sales representative.**

```
LICENSE_MAX_USERS = integer
```

resets (for the current instance) the value of the initialization parameter `LICENSE_MAX_USERS`, which establishes the limit for users connected to your database. Once this limit is reached, more users cannot connect. A value of 0 disables the limit.

**Restriction:** You cannot reduce the limit on users below the current number of users created for the database.

**Do not disable or raise user limits unless you have appropriately upgraded your Oracle license. For more information, contact your Oracle sales representative.**

```
LICENSE_SESSIONS_WARNING = integer
```

resets (for the current instance) the value of the initialization parameter `LICENSE_SESSIONS_WARNING`, which establishes a warning threshold for concurrent usage. Once this threshold is reached, Oracle writes warning messages to the database ALERT file for each subsequent session. Also, users with `RESTRICTED SESSION` system privilege receive warning messages when they begin subsequent sessions. A value of 0 disables the warning threshold.

If you reduce the warning threshold for sessions below the current number of sessions, Oracle writes a message to the ALERT file for all subsequent sessions.

```
LOG_ARCHIVE_DEST = string
```

specifies a valid operating system pathname as the primary destination for all archive redo log file groups.

**Restrictions:** If you set a value for this parameter:

- You cannot have a value for `LOG_ARCHIVE_DEST_n` in your initialization parameter file, nor can you set a value for that parameter using the `ALTER SESSION` or `ALTER SYSTEM` statement.
- You cannot set a value for the parameter `LOG_ARCHIVE_MIN_SUCCEED_DEST` using the `ALTER SESSION` statement.

```
LOG_ARCHIVE_DEST_n = {null_string  
    | {LOCATION=local_pathname | SERVICE=tnsnames_service}  
    [MANDATORY | OPTIONAL] [REOPEN[=integer]]}
```

specifies up to five valid operating system pathnames or Oracle service names (plus other related options) as destinations for archive redo log file groups ( $n =$  integers 1 through 5). For a description of the options, refer to *Oracle8i Reference*.

---

**Restrictions:** If you set a value for this parameter:

- You cannot have definitions for the parameters LOG\_ARCHIVE\_DEST or LOG\_ARCHIVE\_DUPLEX\_DEST in your initialization parameter file, nor can you set values for those parameters using the ALTER SYSTEM statement.
- You cannot start archiving to a specific location using the ALTER SYSTEM ARCHIVE LOG TO *location* statement.

LOG\_ARCHIVE\_DEST\_STATE\_ *n* = {ENABLE | DEFER}

specifies the state associated with the corresponding LOG\_ARCHIVE\_DEST\_ *n* parameter.

- ENABLE specifies that any associated valid destination can be used for archiving. This is the default.
- DEFER specifies that Oracle will not consider for archiving any destination associated with the corresponding LOG\_ARCHIVE\_DEST\_ *n* parameter.

LOG\_ARCHIVE\_DUPLEX\_DEST = string

specifies a valid operating system pathname as the secondary destination for all archive redo log file groups.

**Restriction:** If you set a value for this parameter:

- You must have a definition for LOG\_ARCHIVE\_DEST.
- You cannot have a value for the parameter LOG\_ARCHIVE\_DEST\_ *n* in your initialization parameter file, nor can you set a value for that parameter using the ALTER SYSTEM or ALTER SESSION statement.
- You cannot set a value for the parameter LOG\_ARCHIVE\_MIN\_SUCCEED\_DEST using the ALTER SESSION statement.

LOG\_ARCHIVE\_MAX\_PROCESSES = integer

specifies the number of archiver processes that are invoked. Permitted values are integers 1 through 10, inclusive. The default is 1.

LOG\_ARCHIVE\_MIN\_SUCCEED\_DEST = integer

specifies the minimum number of destinations that must succeed in order for the online log file to be available for reuse.

LOG\_ARCHIVE\_TRACE = integer

controls the type of output information generated by archivelog processes.

**See Also:**

- *Oracle8i Backup and Recovery Guide* for more information on using this parameter.
- *Oracle8i Reference* for a listing of valid values.

LOG\_CHECKPOINT\_INTERVAL = integer

limits to *integer* the number of redo blocks that can exist between an incremental checkpoint and the last block written to the redo log.

```
LOG_CHECKPOINT_TIMEOUT = integer
```

limits the incremental checkpoint to be at the position where the last write to the redo log (sometimes called the "tail of the log") was *integer* seconds ago, and signifies that no buffer will remain dirty (in the cache) for more than *integer* seconds. The default is 1800 seconds.

```
MAX_DUMP_FILE_SIZE = { size | 'UNLIMITED' } [DEFERRED]
```

specifies the trace dump file size upper limit for all user sessions. Specify the maximum *size* as either a nonnegative integer that represents the number of blocks, or as 'UNLIMITED'. If you specify 'UNLIMITED', no upper limit is imposed.

**Multi-Threaded Server Parameters:** When you start your instance, Oracle creates shared server processes and dispatcher processes for the multi-threaded server architecture based on the values of the MTS\_SERVERS and MTS\_DISPATCHERS initialization parameters. You can set the MTS\_SERVERS and MTS\_DISPATCHERS session parameters to perform one of the following operations while the instance is running:

- Create additional shared server processes by increasing the minimum number of shared server processes.
- Terminate existing shared server processes after their current calls finish processing.
- Create more dispatcher processes for a specific protocol, up to a maximum across all protocols specified by the initialization parameter MTS\_MAX\_DISPATCHERS.
- Terminate existing dispatcher processes for a specific protocol after their current user processes disconnect from the instance.

**See Also:**

- *Oracle8i Concepts*
- *Oracle8i Designing and Tuning for Performance*
- *Oracle8i Parallel Server Concepts*.

```
MTS_DISPATCHERS = 'dispatch_clause'
```

```
dispatch_clause:=  
    (PROTOCOL = protocol) |  
    ( ADDRESS = address) |  
    (DESCRIPTION = description )  
    [options_clause]
```

**options\_clause::=**

```
(DISPATCHERS = integer |
SESSIONS = integer |
CONNECTIONS = integer |
TICKS = seconds |
POOL = { 1 | ON | YES | TRUE | BOTH | ({IN|OUT} = ticks) |
0 | OFF | NO | FALSE | ticks} |
MULTIPLEX = {1 | ON | YES | TRUE | 0 | OFF | NO |
FALSE | BOTH | IN | OUT} |
LISTENER = tnsname |
SERVICE = service |
INDEX = integer)
```

modifies or creates the configuration of dispatcher processes. A description of the parameters appears in *Oracle8i Reference*.

You can specify multiple MTS\_DISPATCHERS parameters in a single statement for multiple network protocols.

**See Also:** *Oracle8i Administrator's Guide* for more information on this parameter, see *Net8 Administrator's Guide*.

MTS\_SERVERS = integer

specifies a new minimum number of shared server processes.

OBJECT\_CACHE\_MAX\_SIZE\_PERCENT = integer DEFERRED

specifies the percentage of the optimal cache size that the session object cache can grow past the optimal size.

OBJECT\_CACHE\_OPTIMAL\_SIZE = integer DEFERRED

specifies (in kilobytes) the size to which the session object cache is reduced if it exceeds the maximum size.

OPTIMIZER\_MAX\_PERMUTATIONS = integer NOOVERRIDE

lets you limit the amount of work the optimizer expends on optimizing queries with large joins. The value of *integer* is the number of permutations of the tables the optimizer will consider with large joins.

NOOVERRIDE specifies that this system setting will not override the setting for any session in which this parameter was explicitly set.

PARALLEL\_ADAPTIVE\_MULTI\_USER = {TRUE | FALSE}

specifies that Oracle should vary the degree of parallelism based on the total perceived load on the system.

PARALLEL\_INSTANCE\_GROUP = 'text'

specifies the name of the Oracle Parallel Server instance group to be used for spawning parallel query slaves.

`PARALLEL_THREADS_PER_CPU = integer`

used to compute the degree of parallelism for parallel operations where the degree of parallelism is unset. The default is operating system dependent.

`PLSQL_V2_COMPATIBILITY = { TRUE | FALSE } [DEFERRED]`

modifies the compile-time behavior of PL/SQL programs to allow language constructs that are illegal in Oracle8 and Oracle8i (PL/SQL V3), but were legal in Oracle7 (PL/SQL V2). See *PL/SQL User's Guide and Reference* and *Oracle8i Reference* for more information about this system parameter.

`TRUE` enables Oracle8i PL/SQL V3 programs to execute Oracle7 PL/SQL V2 constructs.

`FALSE` disallows illegal Oracle7 PL/SQL V2 constructs. This is the default.

`QUERY_REWRITE_ENABLED = { TRUE | FALSE } [DEFERRED | NOOVERRIDE]`

enables or disables query rewrite on all materialized views that have not been explicitly disabled. By default, `TRUE` enables query rewrite for all sessions immediately. Query rewrite is superseded and disabled by rule-based optimization (that is, if the `OPTIMIZER_MODE` parameter is set to `RULE`). Also enables or disables use of any function-based indexes defined on the materialized view.

**See Also:** *Oracle8i Data Warehousing Guide* for more information on query rewrite.

- `DEFERRED` specifies that query rewrite is enabled or disabled only for future sessions.
- `NOOVERRIDE` specifies that query rewrite is enabled or disabled for all sessions that have not explicitly set this parameter using `ALTER SESSION`.
- Enabling or disabling query rewrite does not affect queries that have already been compiled, even if they are reissued.
- Enabling or disabling query rewrite does not affect descending indexes.
- A `TRUE` setting has no effect on materialized views that cannot be created with the `ENABLE QUERY REWRITE` clause, such as materialized views created totally or in part from a view.

`QUERY_REWRITE_INTEGRITY = { ENFORCED | TRUSTED | STALE_TOLERATED }`

sets the minimum consistency level for query rewrite for the duration of the instance. The following values are permitted:



- ENFORCED is the safest level. It relies only on system-enforced relationships so that data integrity and correctness can be guaranteed. This level ensures that query rewrite will not use any function-based index or any materialized view that includes a call to a user-defined function.

In addition, this level ensures that query rewrite will not use any dimensional information or any constraints enabled with the RELY keyword.

- TRUSTED specifies that materialized views created with the ON PREBUILT TABLE clause are supported, and trusted but unenforced join relationships are accepted. Query rewrite uses join information from dimensions and enables unenforced constraints with the RELY keyword.
- STALE\_TOLERATED specifies that any stale, usable materialized view may be used.

This parameter does not affect descending indexes.

**See Also:**

- *Oracle8i Data Warehousing Guide* for more information on query rewrite integrity level.
- "CREATE DIMENSION" on page 7-277 for information on dimensions.
- "constraint\_clause" on page 7-233 for information on constraints enabled with the RELY keyword.

REMOTE\_DEPENDENCIES\_MODE = {TIMESTAMP | SIGNATURE}

specifies how dependencies of remote stored procedures are handled by the server.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals*.

RESOURCE\_LIMIT = {TRUE | FALSE}

When you start an instance, Oracle enforces resource limits assigned to users based on the value of the RESOURCE\_LIMIT initialization parameter. This system parameter enables or disables resource limits for subsequent sessions. TRUE enables resource limits. FALSE disables resource limits.

Enabling resource limits only causes Oracle to enforce the resource limits already assigned to users. To choose resource limit values for a user, you must create a profile and assign that profile to the user.

**See Also:** "CREATE PROFILE" on page 7-359 and "CREATE USER" on page 7-451.

RESOURCE\_MANAGER\_PLAN = plan\_name

specifies the name of the resource plan Oracle should use to allocate system resources among resource consumer groups. For information on resource consumer groups and resource plans, refer to *Oracle8i Administrator's Guide*.

## ALTER SYSTEM

---

`SORT_AREA_RETAINED_SIZE = integer DEFERRED`

specifies (in bytes) the maximum amount of memory that each sort operation will retain after the first fetch is done, until the cursor ends. If you do not explicitly set this parameter in the initialization parameter file or dynamically, Oracle uses the value of the `SORT_AREA_SIZE` parameter.

`SORT_AREA_SIZE = integer DEFERRED`

specifies (in bytes) the maximum amount of memory to use for each sort operation. The default is operating system dependent.

`SORT_MULTIBLOCK_READ_COUNT = integer DEFERRED`

specifies the number of database blocks to read each time a sort performs a read from temporary segments. The default is 2.

`STANDBY_ARCHIVE_DEST = string`

specifies a valid operating system pathname as the standby database destination for the archive redo log files.

`TIMED_STATISTICS = {TRUE | FALSE}`

specifies whether the server requests the time from the operating system when generating time-related statistics. The default is `FALSE`.

`TIMED_OS_STATISTICS = integer`

specifies that operating system statistics will be collected when a request is made from a client to the server or when a request completes.

`TRANSACTION_AUDITING = {TRUE | FALSE} DEFERRED`

specifies whether the transaction layer generates a special redo record containing session and user information.

`USE_STORED_OUTLINES = { TRUE | FALSE | 'category_name' } [NOOVERRIDE]`

determines whether the optimizer will use stored outlines to generate execution plans. `USE_STORED_OUTLINES` is not an initialization parameter.

- `TRUE` causes the optimizer to use outlines stored in the `DEFAULT` category when compiling requests.
- `FALSE` specifies that the optimizer should not use stored outlines. This is the default.
- `category_name` causes the optimizer to use outlines stored in the `category_name` category when compiling requests.
- `NOOVERRIDE` specifies that this system setting will not override the setting for any session in which this parameter was explicitly set. If you do not specify `NOOVERRIDE`, this setting takes effect in all sessions.

---

```
USER_DUMP_DEST = 'directory_name'
```

specifies the pathname where Oracle will write debugging trace files on behalf of a user process.

---

## Examples

**Archive Log Examples** The following statement manually archives the redo log file group with the log sequence number 4 in thread number 3:

```
ALTER SYSTEM ARCHIVE LOG THREAD 3 SEQUENCE 4;
```

The following statement manually archives the redo log file group containing the redo log entry with the SCN 9356083:

```
ALTER SYSTEM ARCHIVE LOG CHANGE 9356083;
```

The following statement manually archives the redo log file group containing a member named 'DISKL:LOG6.LOG' to an archived redo log file in the location 'DISKA:[ARCH\$]':

```
ALTER SYSTEM ARCHIVE LOG
  LOGFILE 'diskl:log6.log'
  TO 'diska:[arch$]';
```

**Query Rewrite Example** This statement enables query rewrite in all sessions for all materialized views that have not been explicitly disabled:

```
ALTER SYSTEM SET QUERY_REWRITE_ENABLED = TRUE;
```

**Restricted Session Example** You may want to restrict logons if you are performing application maintenance and you want only application developers with RESTRICTED SESSION system privilege to log on. To restrict logons, issue the following statement:

```
ALTER SYSTEM
  ENABLE RESTRICTED SESSION;
```

You can then terminate any existing sessions using the KILL SESSION clause of the ALTER SYSTEM statement.

After performing maintenance on your application, issue the following statement to allow any user with CREATE SESSION system privilege to log on:

```
ALTER SYSTEM
```

```
DISABLE RESTRICTED SESSION;
```

**Shared Pool Example** You might want to clear the shared pool before beginning performance analysis. To clear the shared pool, issue the following statement:

```
ALTER SYSTEM FLUSH SHARED_POOL;
```

**CHECKPOINT Example** The following statement forces a checkpoint:

```
ALTER SYSTEM CHECKPOINT;
```

**Resource Limit Example** This ALTER SYSTEM statement dynamically enables resource limits:

```
ALTER SYSTEM SET RESOURCE_LIMIT = TRUE;
```

**Multi-Threaded Server Examples** The following statement changes the minimum number of shared server processes to 25:

```
ALTER SYSTEM SET MTS_SERVERS = 25;
```

If there are currently fewer than 25 shared server processes, Oracle creates more. If there are currently more than 25, Oracle terminates some of them when they are finished processing their current calls if the load could be managed by the remaining 25.

The following statement dynamically changes the number of dispatcher processes for the TCP/IP protocol to 5 and the number of dispatcher processes for the DECNet protocol to 10:

```
ALTER SYSTEM
  SET MTS_DISPATCHERS =
    '( INDEX=0 ) ( PROTOCOL=TCP ) ( DISPATCHERS=5 ) ',
    '( INDEX=1 ) ( PROTOCOL=DECNet ) ( DISPATCHERS=10 ) ' ;
```

If there are currently fewer than 5 dispatcher processes for TCP, Oracle creates new ones. If there are currently more than 5, Oracle terminates some of them after the connected users disconnect.

If there are currently fewer than 10 dispatcher processes for DECNet, Oracle creates new ones. If there are currently more than 10, Oracle terminates some of them after the connected users disconnect.

If there are currently existing dispatchers for another protocol, the above statement does not affect the number of dispatchers for that protocol.

**Licensing Examples** The following statement dynamically changes the limit on sessions for your instance to 64 and the warning threshold for sessions on your instance to 54:

```
ALTER SYSTEM
  SET LICENSE_MAX_SESSIONS = 64
  LICENSE_SESSIONS_WARNING = 54;
```

If the number of sessions reaches 54, Oracle writes a warning message to the ALERT file for each subsequent session. Also, users with `RESTRICTED SESSION` system privilege receive warning messages when they begin subsequent sessions.

If the number of sessions reaches 64, only users with `RESTRICTED SESSION` system privilege can begin new sessions until the number of sessions falls below 64 again.

The following statement dynamically disables the limit for sessions on your instance. After you issue the above statement, Oracle no longer limits the number of sessions on your instance.

```
ALTER SYSTEM SET LICENSE_MAX_SESSIONS = 0;
```

The following statement dynamically changes the limit on the number of users in the database to 200. After you issue the above statement, Oracle prevents the number of users in the database from exceeding 200.

```
ALTER SYSTEM SET LICENSE_MAX_USERS = 200;
```

**SWITCH LOGFILE Example** You may want to force a log switch to drop or rename the current redo log file group or one of its members, because you cannot drop or rename a file while Oracle is writing to it. The forced log switch affects only your instance's redo log thread. The following statement forces a log switch:

```
ALTER SYSTEM
  SWITCH LOGFILE;
```

**Distributed Recovery Example** The following statement enables distributed recovery:

```
ALTER SYSTEM ENABLE DISTRIBUTED RECOVERY;
```

You may want to disable distributed recovery for demonstration or testing purposes. You can disable distributed recovery in both single-process and multiprocess mode with the following statement:

```
ALTER SYSTEM DISABLE DISTRIBUTED RECOVERY;
```

When your demonstration or testing are complete, you can then enable distributed recovery again by issuing an `ALTER SYSTEM` statement with the `ENABLE DISTRIBUTED RECOVERY` clause.

**KILL SESSION Example** You may want to kill the session of a user that is holding resources needed by other users. The user receives an error message indicating that the session has been killed. That user can no longer make calls to the database without beginning a new session. Consider this data from the `V$SESSION` dynamic performance table:

```
SELECT sid, serial, username
FROM v$session
```

SID	SERIAL	USERNAME
1	1	
2	1	
3	1	
4	1	
5	1	
7	1	
8	28	OPS\$BQUIGLEY
10	211	OPS\$SWIFT
11	39	OPS\$OBRIEN
12	13	SYSTEM
13	8	SCOTT

The following statement kills the session of the user `SCOTT` using the `SID` and `SERIAL#` values from `V$SESSION`:

```
ALTER SYSTEM KILL SESSION '13, 8';
```


**DISCONNECT SESSION Example** The following statement disconnects user `SCOTT`'s session, using the `SID` and `SERIAL#` values from `V$SESSION`:

```
ALTER SYSTEM DISCONNECT SESSION '13, 8' POST_TRANSACTION;
```


**See Also:** *Oracle8i Parallel Server Concepts* and *Oracle8i Designing and Tuning for Performance* for more information about application failover.

# ALTER TABLE

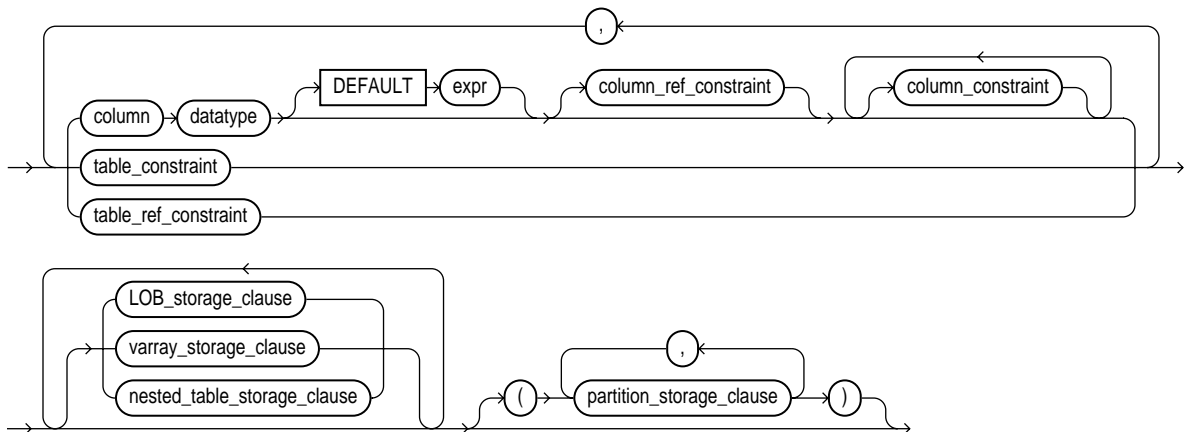
## Syntax



## ALTER TABLE




### add\_column\_options::=

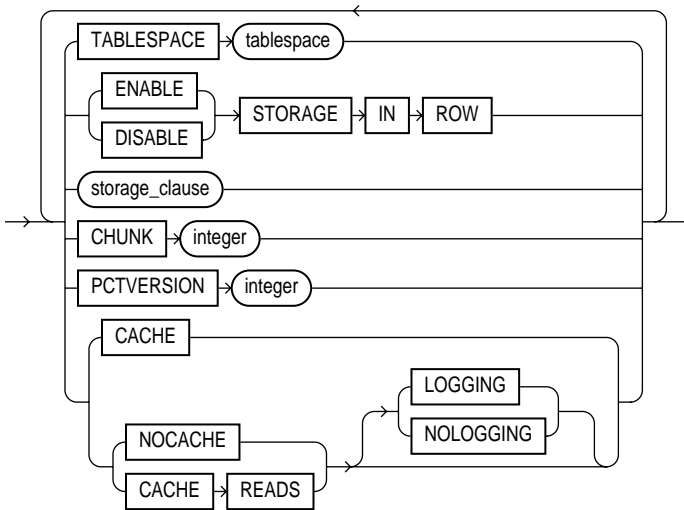


**column\_constraint, table\_constraint, column\_ref\_constraint, table\_ref\_constraint, constraint\_state:** See the "constraint\_clause" on page 7-233.


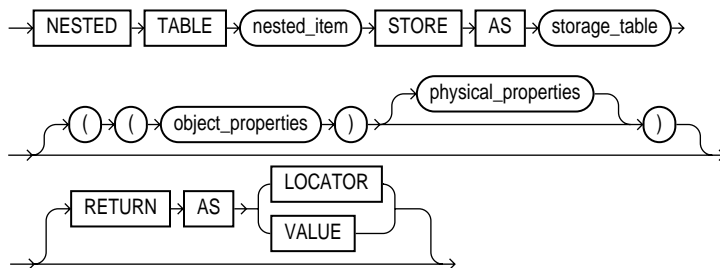
### LOB\_storage\_clause::=






**LOB\_parameters::=**

**storage\_clause**: See "storage\_clause" on page 7-605.

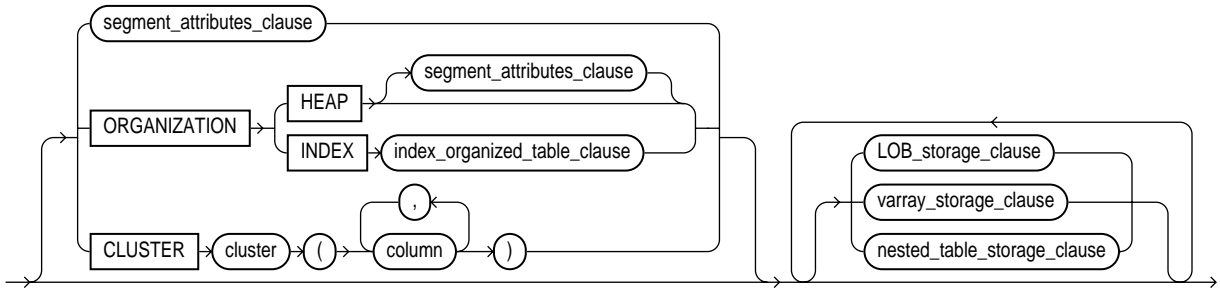
**varray\_storage\_clause::=****nested\_table\_storage\_clause::=**

# ALTER TABLE


## object\_properties::=




## physical\_properties::=




## partition\_storage\_clause::=




## modify\_column\_options::=




**move\_table\_clause::=**




**segment\_attributes\_clause::=**



**physical\_attributes\_clause::=**




**index\_organized\_table\_clause::=**



## ALTER TABLE

---


### compression\_clause::=




### index\_organized\_overflow\_clause::=



### modify\_collection\_retrieval\_clause::=



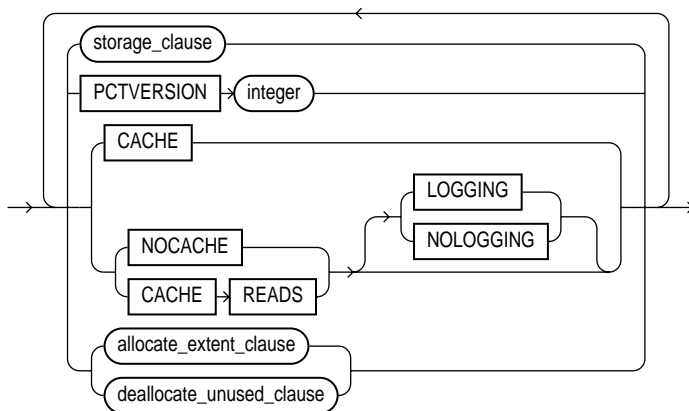
### modify\_storage\_clauses::=



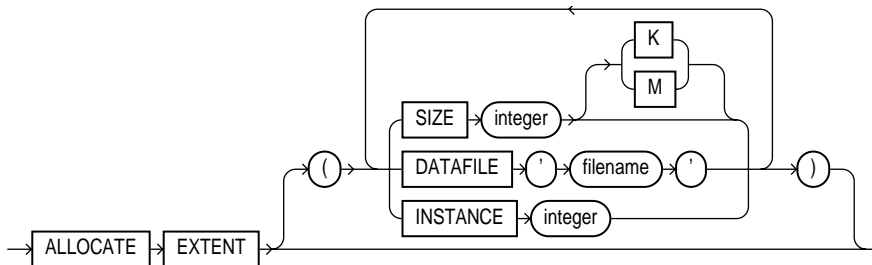
### modify\_LOB\_storage\_clause::=




**modify\_LOB\_storage\_parameters::=**




**allocate\_extent\_clause::=**



**deallocate\_unused\_clause::=**




**modify\_varray\_storage\_clause::=**

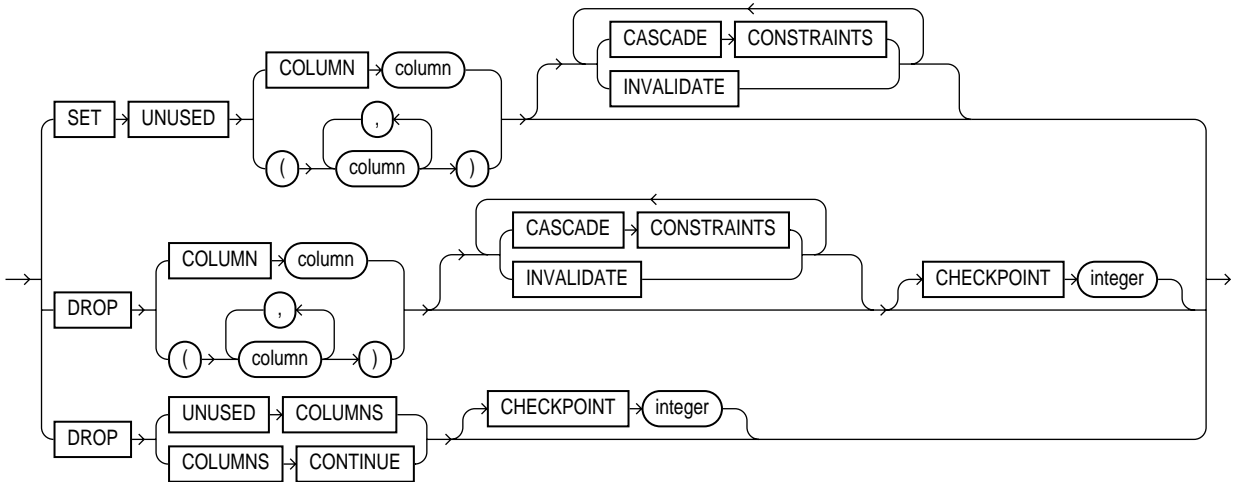


## ALTER TABLE


**drop\_constraint\_clause::=**






**drop\_column\_clause::=**



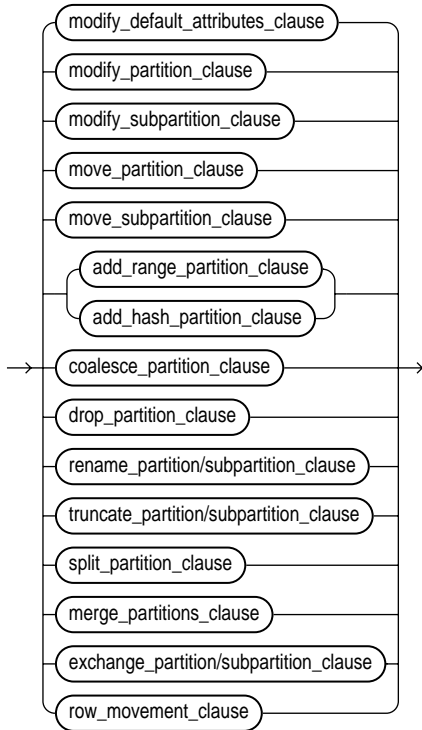
**records\_per\_block\_clause::=**



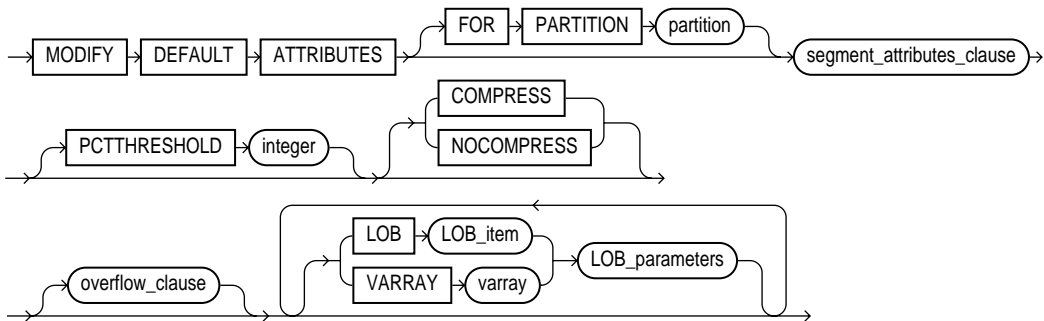
**alter\_overflow\_clause::=****overflow\_clause::=****add\_overflow\_clause::=**

## ALTER TABLE


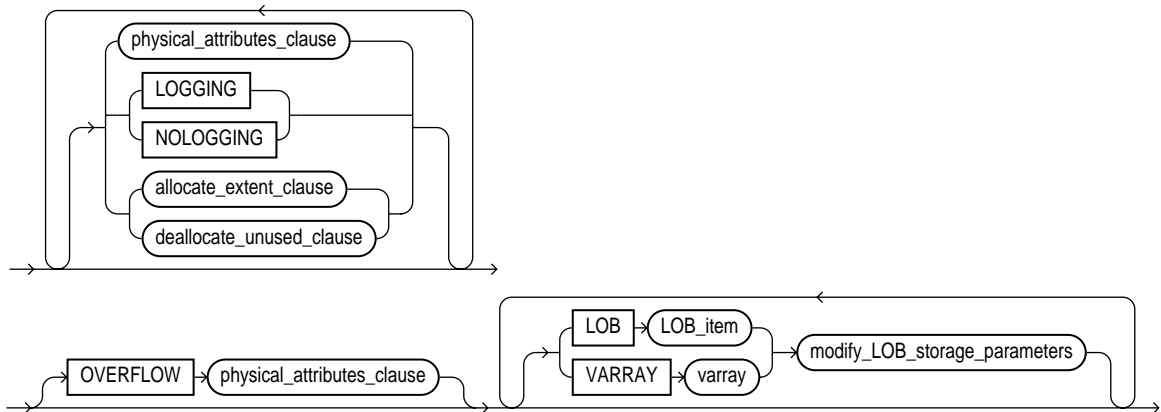


### partitioning\_clauses::=



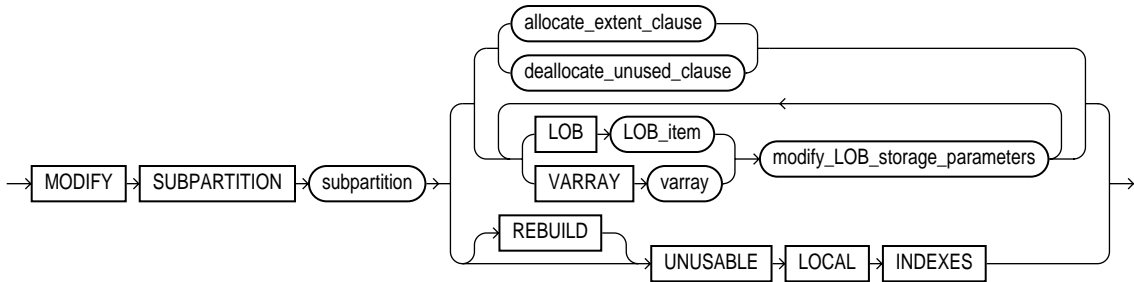
### modify\_default\_attributes\_clause::=






**modify\_partition\_clause::=****partition\_attributes::=****add\_subpartition\_clause::=****subpartition\_description::=**

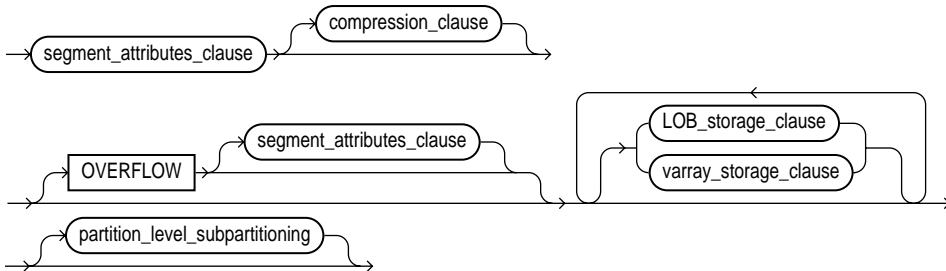
**modify\_subpartition\_clause::=**




**move\_partition\_clause::=**







**partition\_description::=**




**partition\_level\_subpartitioning::=**




**hash\_partitioning\_storage\_clause::=****move\_subpartition\_clause::=****add\_range\_partition\_clause::=****add\_hash\_partition\_clause::=****coalesce\_partition\_clause::=****drop\_partition\_clause::=**


**rename\_partition/ subpartition\_clause::=**




**truncate\_partition\_clause/truncate\_subpartition\_clause::=**

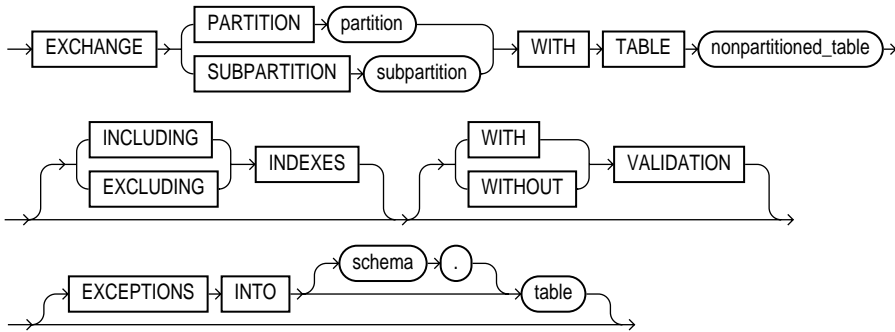


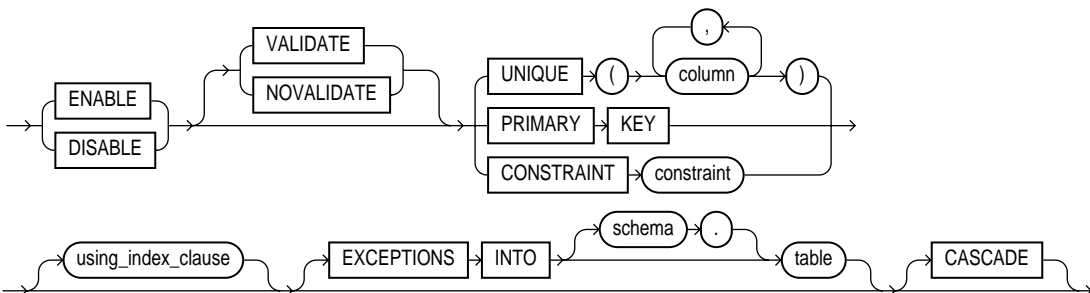


**split\_partition\_clause::=**

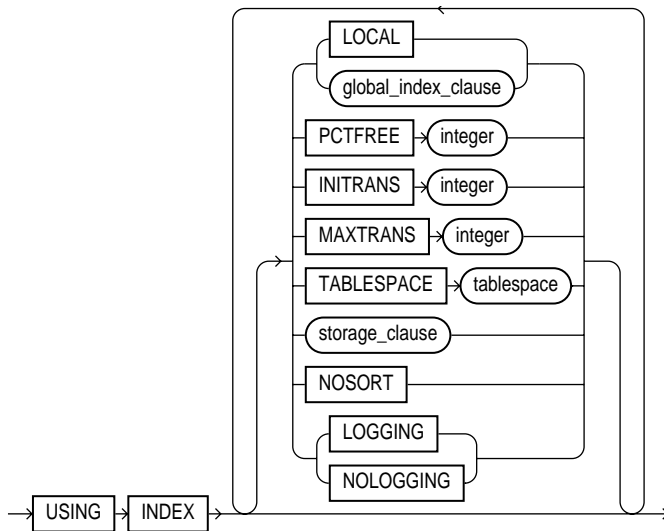


**merge\_partitions\_clause::=**



**exchange\_partition\_clause/exchange\_subpartition\_clause::=****row\_movement\_clause::=****parallel\_clause::=****enable\_disable\_clause::=**

**using\_index\_clause::=**



## Purpose

To alter the definition of a nonpartitioned table, a partitioned table, a table partition, or a table subpartition.

## Prerequisites

The table must be in your own schema, or you must have `ALTER` privilege on the table, or you must have `ALTER ANY TABLE` system privilege. For some operations you may also need the `CREATE ANY INDEX` privilege.

In addition, if you are not the owner of the table, you need the `DROP ANY TABLE` privilege in order to use the `drop_partition_clause` or `truncate_partition_clause`.

You must also have space quota in the tablespace in which space is to be acquired in order to use the `add_partition_clause`, `modify_partition_clause`, `move_partition_clause`, and `split_partition_clause`.

To enable a `UNIQUE` or `PRIMARY KEY` constraint, you must have the privileges necessary to create an index on the table. You need these privileges because Oracle creates an index on the columns of the unique or primary key in the schema containing the table. See "[CREATE INDEX](#)" on page 7-291.

To enable or disable triggers, the triggers must be in your schema or you must have the ALTER ANY TRIGGER system privilege.

To use an object type in a column definition when modifying a table, either that object must belong to the same schema as the table being altered, or you must have either the EXECUTE ANY TYPE system privilege or the EXECUTE schema object privilege for the object type.

## Keywords and Parameters

---

The clauses described below have specialized meaning in the ALTER TABLE statement. For descriptions of the remaining keywords, see "CREATE TABLE" on page 4-381.

---

**Note:** Operations performed by the ALTER TABLE statement can cause Oracle to invalidate procedures and stored functions that access the table. For information on how and when Oracle invalidates such objects, see *Oracle8i Concepts*.

---

<i>schema</i>	is the schema containing the table. If you omit <i>schema</i> , Oracle assumes the table is in your own schema.
<i>table</i>	is the name of the table to be altered.  You can modify, or drop columns from, or rename a temporary table. However, for a temporary table, you cannot: <ul style="list-style-type: none"> <li>■ Add columns of nested-table or varray type. You can add columns of other types.</li> <li>■ Specify referential integrity (foreign key) constraints for an added or modified column</li> <li>■ Specify the following clauses of the <i>LOB_storage_clause</i> for an added or modified LOB column: <i>TABLESPACE</i>, <i>storage_clause</i>, <i>LOGGING</i>   <i>NOLOGGING</i>, or the <i>LOB_index_clause</i>.</li> <li>■ Specify the <i>physical_attribute_clause</i>, <i>nested_table_storage_clause</i>, <i>parallel_clause</i>, <i>allocate_extent_clause</i>, <i>deallocate_unused_clause</i>, or any of the <i>index-organized table clauses</i></li> <li>■ Exchange partitions between a partition and a temporary table</li> <li>■ Specify <i>LOGGING</i> or <i>NOLOGGING</i></li> <li>■ Specify <i>MOVE</i></li> </ul>

---

**Note:** If you alter a table that is a master table for one or more materialized views, the materialized views are marked *INVALID*. Invalid materialized views cannot be used by query rewrite and cannot be refreshed. To revalidate a materialized view, see "ALTER MATERIALIZED VIEW / SNAPSHOT" on page 7-47.

**See Also:** *Oracle8i Data Warehousing Guide* for more information on materialized views in general.

<b>ADD</b>	adds a column or integrity constraint. For a description of the keywords and parameters of this clause, see " <a href="#">CREATE TABLE</a> " on page 4-381.
<i>add_column_options</i>	<p>If you add a column, the initial value of each row for the new column is null unless you specify the <code>DEFAULT</code> clause. In this case, Oracle updates each row in the new column with the value you specify for <code>DEFAULT</code>. This update operation, in turn, fires any <code>AFTER UPDATE</code> triggers defined on the table.</p> <p>You can add an overflow data segment to each partition of a partitioned index-organized table.</p> <p>You can add LOB columns to nonpartitioned and partitioned tables. You can specify LOB storage at the table and at the partition or subpartition level.</p> <p>If you previously created a view with a query that used the "SELECT *" syntax to select all columns from <i>table</i>, and you now add a column to <i>table</i>, Oracle does not automatically add the new column to the view. To add the new column to the view, re-create the view using the <code>CREATE VIEW</code> statement with the <code>OR REPLACE</code> clause. See "<a href="#">CREATE VIEW</a>" on page 7-456.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot add a LOB column to a partitioned index-organized table. (This restriction does not apply to nonpartitioned index-organized tables.)</li> <li>■ You cannot add a column with a <code>NOT NULL</code> constraint if <i>table</i> has any rows unless you also specify the <code>DEFAULT</code> clause.</li> <li>■ If you specify this clause for an index-organized table, you cannot specify any other clauses in the same statement.</li> </ul>
<b>DEFAULT</b>	<p>specifies a default for a new column or a new default for an existing column. Oracle assigns this value to the column if a subsequent <code>INSERT</code> statement omits a value for the column. If you are adding a new column to the table and specify the default value, Oracle inserts the default column value into all rows of the table.</p> <p>The datatype of the default value must match the datatype specified for the column. The column must also be long enough to hold the default value. A <code>DEFAULT</code> expression cannot contain references to other columns, the pseudocolumns <code>CURRVAL</code>, <code>NEXTVAL</code>, <code>LEVEL</code>, and <code>ROWNUM</code>, or date constants that are not fully specified.</p>
<i>table_ref_constraint</i>	<p>These clauses let you further describe a column of type <code>REF</code>. The only difference between these clauses is that you specify <i>table_ref</i> from the table level, so you must identify the <code>REF</code> column or attribute you are defining. You specify <i>column_ref</i> after you have already identified the <code>REF</code> column or attribute.</p> <p>For syntax and description of these constraints, including restrictions, see the "<a href="#">constraint_clause</a>" on page 7-233.</p>
<i>column_ref_constraint</i>	



---

<i>column_constraint</i>	adds or removes a NOT NULL constraint to or from an existing column. You cannot use this clause to modify any other type of constraint using ALTER TABLE. See the " <a href="#">constraint_clause</a> " on page 7-233.
<i>table_constraint</i>	adds or modifies an integrity constraint on the table. See the " <a href="#">constraint_clause</a> " on page 7-233.
<i>LOB_storage_clause</i>	specifies the LOB storage characteristics for the newly added LOB column. You cannot use this clause to modify an existing LOB column. Instead, you must use the <i>modify_LOB_storage_clause</i> .
	<b>Restrictions:</b>
	<ul style="list-style-type: none"> <li>■ The only parameter of <i>LOB_parameters</i> you can specify for a hash partition or hash subpartition is TABLESPACE.</li> <li>■ You cannot specify the <i>LOB_index_clause</i> if <i>table</i> is partitioned.</li> </ul>
<i>lob_item</i>	is the LOB column name or LOB object attribute for which you are explicitly defining tablespace and storage characteristics that are different from those of the table.
<i>lob_segname</i>	specifies the name of the LOB data segment. You cannot use <i>lob_segname</i> if more than one <i>lob_item</i> is specified.
ENABLE   DISABLE STORAGE IN ROW	<p>specifies whether the LOB value is stored in the row (inline) or outside of the row. (The LOB locator is always stored in the row regardless of where the LOB value is stored.)</p> <ul style="list-style-type: none"> <li>■ ENABLE specifies that the LOB value is stored inline if its length is less than approximately 4000 bytes minus system control information. This is the default.</li> <li>■ DISABLE specifies that the LOB value is stored outside of the row regardless of the length of the LOB value.</li> </ul> <p><b>Restriction:</b> You cannot change STORAGE IN ROW once it is set. Therefore, you can specify this clause only as part of the <i>add_column_options</i> clause, not as part of the <i>modify_column_options</i> clause.</p>
<i>CHUNK integer</i>	<p>specifies the number of bytes to be allocated for LOB manipulation. If <i>integer</i> is not a multiple of the database block size, Oracle rounds up (in bytes) to the next multiple. For example, if the database block size is 2048 and <i>integer</i> is 2050, Oracle allocates 4096 bytes (2 blocks). The maximum value is 32768 (32 K), which is the largest Oracle block size allowed. The default <i>CHUNK</i> size is one Oracle database block.</p> <p>You cannot change the value of <i>CHUNK</i> once it is set.</p> <hr/> <p><b>Note:</b> The value of <i>CHUNK</i> must be less than or equal to the value of <i>NEXT</i> (either the default value or that specified in the storage clause). If <i>CHUNK</i> exceeds the value of <i>NEXT</i>, Oracle returns an error.</p>

---

<i>PCTVERSION integer</i>	is the maximum percentage of overall LOB storage space used for creating new versions of the LOB. The default value is 10, meaning that older versions of the LOB data are not overwritten until 10% of the overall LOB storage space is used.
<i>LOB_index_clause</i>	<p>This clause is deprecated as of Oracle8i. Oracle generates an index for each LOB column. The LOB indexes are system named and system managed, and reside in the same tablespace as the LOB data segments.</p> <p>It is still possible for you to specify this clause in some cases. However, Oracle Corporation strongly recommends that you no longer do so.</p> <p>For information on how Oracle manages LOB indexes in tables migrated from earlier versions, see <i>Oracle8i Migration</i>.</p>
<i>varray_storage_clause</i>	<p>lets you specify separate storage characteristics for the LOB in which a varray will be stored. In addition, if you specify this clause, Oracle will always store the varray in a LOB, even if it is small enough to be stored inline.</p> <p><b>Restriction:</b> You cannot specify the <i>TABLESPACE</i> clause of <i>LOB_parameters</i> as part of this clause. The LOB tablespace for a varray defaults to the containing table's tablespace.</p>
<i>nested_table_storage_clause</i>	<p>enables you to specify separate storage characteristics for a nested table, which in turn enables you to define the nested table as an index-organized table. You must include this clause when creating a table with columns or column attributes whose type is a nested table. (Clauses within this clause that function the same way they function for parent object tables are not repeated here.)</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify the <i>parallel_clause</i>.</li> <li>■ You cannot specify <i>TABLESPACE</i> (as part of the <i>segment_attributes_clause</i>) for a nested table. The tablespace is always that of the parent table.</li> </ul>
<i>nested_item</i>	is the name of a column (or a top-level attribute of the table's object type) whose type is a nested table.
<i>storage_table</i>	is the name of the table where the rows of <i>nested_item</i> reside. The storage table is created in the same schema and the same tablespace as the parent table.
<i>partition_storage_clause</i>	<p>lets you specify a separate <i>LOB_storage_clause</i> or <i>varray_storage_clause</i> for each partition. You must specify the partitions in the order of partition position.</p> <p>If you do not specify a <i>LOB_storage_clause</i> or <i>varray_storage_clause</i> for a particular partition, the storage characteristics are those specified for the LOB item at the table level. If you also did not specify any storage characteristics at the table level for the LOB item, Oracle stores the LOB data partition in the same tablespace as the table partition to which it corresponds.</p> <p><b>Restriction:</b> You can specify only one list of <i>partition_storage_clauses</i> per <i>ALTER TABLE</i> statement, and all <i>LOB_storage_clauses</i> and <i>varray_storage_clauses</i> must precede the list of <i>partition_storage_clauses</i>.</p>

---

**MODIFY** modifies the definition of an existing column. If you omit any of the optional parts of the column definition (datatype, default value, or column constraint), these parts remain unchanged.

*modify\_column\_options*

- You can change a CHAR column to VARCHAR2 (or VARCHAR) and a VARCHAR2 (or VARCHAR) to CHAR only if the column contains nulls in all rows or if you do not attempt to change the column size.
- You can change any column's datatype or decrease any column's size if all rows for the column contain nulls.
- You can always increase the size of a character or raw column or the precision of a numeric column, whether or not all the columns contain nulls.

**Restrictions:**

- You cannot modify the datatype or length of a column that is part of a table or index partitioning or subpartitioning key.
- You cannot modify the definition of a column on which a domain index has been built.
- If you specify this clause for an index-organized table, you cannot specify any other clauses in the same statement.

*column* is the name of the column to be added or modified.

The only type of integrity constraint that you can add to an existing column using the MODIFY clause with the column constraint syntax is a NOT NULL constraint, and only if the column contains no nulls. To define other types of integrity constraints (UNIQUE, PRIMARY KEY, referential integrity, and CHECK constraints) on existing columns, using the ADD clause and the table constraint syntax.

*datatype* specifies a new datatype for an existing column.

You can omit the datatype only if the statement also designates the column as part of the foreign key of a referential integrity constraint. Oracle automatically assigns the column the same datatype as the corresponding column of the referenced key of the referential integrity constraint.

If you change the datatype of a column in a materialized view container table, the corresponding materialized view is invalidated. To revalidate a materialized view, see "[ALTER MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-47.

**Restrictions:**

- You cannot specify a column of datatype ROWID for an index-organized table, but you can specify a column of type UROWID.
- You cannot change a column's datatype to LOB or REF.

---

MODIFY CONSTRAINT <i>constraint</i>	<p>modifies the state of an existing constraint named <i>constraint</i>. For a description of all the keywords and parameters of <i>constraint_state</i>, see the "<a href="#">constraint_clause</a>" on page 7-233.</p>
<i>move_table_clause</i>	<p>For a <b>heap-organized table</b>, use the <i>segment_attributes_clause</i> of the syntax. The <i>move_table_clause</i> lets you relocate data of a nonpartitioned table into a new segment, optionally in a different tablespace, and optionally modify any of its storage attributes.</p> <p>You can also move any LOB data segments associated with the table using the <i>LOB_storage_clause</i>. (LOB items not specified in this clause are not moved.)</p> <p>For an <b>index-organized table</b>, use the <i>index_organized_table_clause</i> of the syntax. The <i>move_table_clause</i> rebuilds the index-organized table's primary key index B*-tree. The overflow data segment is not rebuilt unless the <code>OVERFLOW</code> keyword is explicitly stated, with two exceptions:</p> <ul style="list-style-type: none"> <li>■ If you alter the values of <code>PCTTHRESHOLD</code> or the <code>INCLUDING</code> column as part of this <code>ALTER TABLE</code> statement, the overflow data segment is rebuilt.</li> <li>■ If any of out-of-line columns (LOBs, varrays, nested table columns) in the index-organized table are moved explicitly, then the overflow data segment is also rebuilt.</li> </ul> <p>The index and data segments of LOB columns are not rebuilt unless you specify the LOB columns explicitly as part of this <code>ALTER TABLE</code> statement.</p>
ONLINE	<p>specifies that DML operations on the index-organized table are allowed during rebuilding of the table's primary key index B*-tree.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You can specify this clause only for a nonpartitioned index-organized table.</li> <li>■ Parallel DML is not supported during online <code>MOVE</code>. If you specify <code>ONLINE</code> and then issue parallel DML statements, Oracle returns an error.</li> </ul>
<i>compression_clause</i>	<p>enables and disables key compression in an index-organized table.</p> <ul style="list-style-type: none"> <li>■ <code>COMPRESS</code> enables key compression, which eliminates repeated occurrence of primary key column values in index-organized tables. Use <i>integer</i> to specify the prefix length (number of prefix columns to compress).</li> </ul> <p>The valid range of prefix length values is from 1 to the number of primary key columns minus 1. The default prefix length is the number of primary key columns minus 1.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>- You can specify this clause only for an index-organized table.</li> <li>- You can specify compression for a partition of an index-organized table only if compression has been specified at the table level.</li> </ul>

- NOCOMPRESS disables key compression in index-organized tables. This is the default.

TABLESPACE specifies the tablespace into which the rebuilt index-organized table is stored.

#### Restrictions:

- If you specify MOVE, it must be the first clause. For an index-organized table, the only clauses outside this clause that are allowed are the *physical\_attribute\_clause* and the *parallel\_clause*. For heap-organized tables, you can specify those two clauses and the *LOB\_storage\_clauses*.
- You cannot MOVE an entire partitioned table (either heap or index organized). You must move individual partitions or subpartitions. See "[move\\_partition\\_clause](#)" on page 7-157 and "[move\\_subpartition\\_clause](#)" on page 7-158.

#### Notes regarding LOBs:

For any LOB columns you specify in this clause:

- Oracle drops the old LOB data segment and corresponding index segment and creates new segments, even if you do not specify a new tablespace.
- If the LOB index in *table* resided in a different tablespace from the LOB data, Oracle collocates the LOB index with the LOB data in the LOB data's tablespace after the move.

*physical\_attributes\_clause*

changes the value of PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters and storage characteristics. See the PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters of "[CREATE TABLE](#)" on page 4-381 and the "[storage\\_clause](#)".

**Restriction:** You cannot specify the PCTUSED parameter for the index segment of an index-organized table.

#### WARNING:

- For a nonpartitioned table, the values you specify **override any values specified for the table at create time**.
- For a range- or hash-partitioned table, the values you specify are the default values for the table and the **actual values for every existing partition, overriding any values already set for the partitions**. To change default table attributes without overriding existing partition values, use the *modify\_default\_attributes\_clause*.
- For a composite-partitioned table, the values you specify are the default values for the table and all partitions of the table and the **actual values for all subpartitions of the table, overriding any values already set for the subpartitions**. To change default partition attributes without overriding existing subpartition values, use the *modify\_default\_attributes\_clause* with the FOR PARTITION clause.

*modify\_collection\_retrieval\_clause*

changes what is returned when a collection item is retrieved from the database.

---

<i>collection_item</i>	is the name of a column-qualified attribute whose type is nested table or varray.
RETURN AS	<p>specifies what Oracle returns as the result of a query.</p> <ul style="list-style-type: none"> <li>■ LOCATOR specifies that a unique locator for the nested table is returned.</li> <li>■ VALUE specifies that a copy of the nested table itself is returned.</li> </ul>
<i>modify_storage_clauses:</i>	
<i>modify_LOB_storage_clause</i>	<p>modifies the physical attributes of the LOB <i>lob_item</i>. You can specify only one <i>lob_item</i> for each <i>modify_LOB_storage_clause</i>.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot modify the value of the INITIAL parameter in the <i>storage_clause</i> when modifying the LOB storage attributes.</li> <li>■ You cannot specify both the <i>allocate_extent_clause</i> and the <i>deallocate_unused_clause</i> in the same statement.</li> </ul>
<i>modify_varray_storage_clause</i>	<p>lets you change the storage characteristics of an existing LOB in which a varray is stored.</p> <p><b>Restriction:</b> You cannot specify the TABLESPACE clause of <i>LOB_parameters</i> as part of this clause. The LOB tablespace for a varray defaults to the containing table's tablespace.</p>
<i>drop_constraint_clause</i>	<p>drops an integrity constraint from the database. Oracle stops enforcing the constraint and removes it from the data dictionary. You can specify only one constraint for each <i>drop_constraint_clause</i>, but you can specify multiple <i>drop_constraint_clauses</i> in one statement.</p> <p>PRIMARY KEY      drops the table's PRIMARY KEY constraint.</p> <p>UNIQUE            drops the UNIQUE constraint on the specified columns.</p> <p>CONSTRAINT       drops the integrity constraint named <i>constraint</i>.</p> <p>CASCADE          drops all other integrity constraints that depend on the dropped integrity constraint.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot drop a UNIQUE or PRIMARY KEY constraint that is part of a referential integrity constraint without also dropping the foreign key. To drop the referenced key and the foreign key together, use the CASCADE clause. If you omit CASCADE, Oracle does not drop the PRIMARY KEY or UNIQUE constraint if any foreign key references it.</li> <li>■ You cannot drop a primary key constraint (even with the CASCADE clause) on a table that uses the primary key as its object identifier (OID).</li> <li>■ If you drop a referential integrity constraint on a REF column, the REF column remains scoped to the referenced table.</li> <li>■ You cannot drop the scope of the column.</li> </ul>

---

*drop\_column\_  
clause*

lets you free space in the database by dropping columns you no longer need, or by marking them to be dropped at a future time when the demand on system resources is less.

SET UNUSED

marks one or more columns as unused. Specifying this clause does not actually remove the target columns from each row in the table (that is, it does not restore the disk space used by these columns). Therefore, the response time is faster than it would be if you execute the `DROP` clause.

You can view all tables with columns marked as unused in the data dictionary views `USER_UNUSED_COL_TABS`, `DBA_UNUSED_COL_TABS`, and `ALL_UNUSED_COL_TABS`. For information on these views, see *Oracle8i Reference*.

Unused columns are treated as if they were dropped, even though their column data remains in the table's rows. After a column has been marked as unused, you have no access to that column. A "SELECT \*" query will not retrieve data from unused columns. In addition, the names and types of columns marked unused will not be displayed during a `DESCRIBE`, and you can add to the table a new column with the same name as an unused column.

---

**Note:** Until you actually drop these columns, they continue to count toward the absolute limit of 1000 columns per table. However, as with all DDL statements, you cannot roll back the results of this clause. That is, you cannot issue `SET USED` counterpart to retrieve a column that you have `SET UNUSED`.

Also, if you mark a column of datatype `LONG` as `UNUSED`, you cannot add another `LONG` column to the table until you actually drop the unused `LONG` column.

**See Also:** "[CREATE TABLE](#)" on page 4-381 for more information on the 1000 column limit.

DROP

---

removes the column descriptor and the data associated with the target column from each row in the table. If you explicitly drop a particular column, all columns currently marked as unused in the target table are dropped at the same time.

---

When the column data is dropped:

- All indexes defined on any of the target columns are also dropped.
- All constraints that reference a target column are removed.
- If any statistics types are associated with the target columns, Oracle disassociates the statistics from the column with the `FORCE` option and drops any statistics collected using the statistics type.

**See Also:** "[DISASSOCIATE STATISTICS](#)" on page 7-470 for more information on disassociating statistics types.

---

**Note:** If a constraint also references a nontarget column, Oracle returns an error and does not drop the column unless you have specified the `CASCADE CONSTRAINTS` clause. If you have specified that clause, Oracle removes all constraints that reference any of the target columns.

DROP UNUSED  
COLUMNS

removes from the table all columns currently marked as unused. Use this statement when you want to reclaim the extra disk space from unused columns in the table. If the table contains no unused columns, the statement returns with no errors.

*column*

specifies one or more columns to be set as unused or dropped. Use the `COLUMN` keyword only if you are specifying only one column. If you specify a column list, it cannot contain duplicates.

CASCADE  
CONSTRAINTS

drops all referential integrity constraints that refer to the primary and unique keys defined on the dropped columns, and drops all multicolumn constraints defined on the dropped columns. If any constraint is referenced by columns from other tables or remaining columns in the target table, then you must specify `CASCADE CONSTRAINTS`. Otherwise, the statement aborts and an error is returned.

INVALIDATE

---

**Note:** Currently, Oracle executes this clause regardless of whether you specify the keyword `INVALIDATE`.

Oracle invalidates all dependent objects, such as views, triggers, and stored program units. Object invalidation is a recursive process. Therefore, all directly dependent *and* indirectly dependent objects are invalidated. However, only local dependencies are invalidated, because Oracle manages remote dependencies differently from local dependencies.

An object invalidated by this statement is automatically revalidated when next referenced. You must then correct any errors that exist in that object before referencing it.

**See Also:** *Oracle8i Concepts* for more information on dependencies.



CHECKPOINT	<p>specifies that a checkpoint for the drop column operation will be applied after processing <i>integer</i> rows; <i>integer</i> is optional and must be greater than zero. If <i>integer</i> is greater than the number of rows in the table, Oracle applies a checkpoint after all the rows have been processed. If you do not specify <i>integer</i>, Oracle sets the default of 512.</p> <p>Checkpointing cuts down the amount of undo logs accumulated during the drop column operation to avoid running out of rollback segment space. However, if this statement is interrupted after a checkpoint has been applied, the table remains in an unusable state. While the table is unusable, the only operations allowed on it are DROP TABLE, TRUNCATE TABLE, and ALTER TABLE DROP COLUMNS CONTINUE (described below).</p> <p>You cannot use this clause with SET UNUSED, because that clause does not remove column data.</p>
DROP COLUMNS CONTINUE	<p>continues the drop column operation from the point at which it was interrupted. Submitting this statement while the table is in a valid state results in an error.</p>

**Restrictions on the *drop\_column\_clause*:**

- Each of the parts of this clause can be specified only once in the statement and cannot be mixed with any other ALTER TABLE clauses. For example, the following statements are not allowed:

```
ALTER TABLE t1 DROP COLUMN f1 DROP (f2);
ALTER TABLE t1 DROP COLUMN f1 SET UNUSED (f2);
ALTER TABLE t1 DROP (f1) ADD (f2 NUMBER);
ALTER TABLE t1 SET UNUSED (f3)
    ADD (CONSTRAINT ck1 CHECK (f2 > 0));
```

- You can drop an object type column only as an entity. Dropping an attribute from an object type column is not allowed.
- If you drop a nested table column, its storage table is removed.
- If you drop a LOB column, the LOB data and its corresponding LOB index segment are removed.
- If you drop a BFILE column, only the locators stored in that column are removed, not the files referenced by the locators.
- You can drop a column from an index-organized table only if it is not a primary key column. The primary key constraint of an index-organized table can never be dropped, so you cannot drop a primary key column even if you have specified CASCADE CONSTRAINTS.

- You can export tables with dropped or unused columns. However, you can import a table only if all the columns specified in the export files are present in the table (that is, none of those columns has been dropped or marked unused). Otherwise, Oracle returns an error.
- You cannot drop a column on which a domain index has been built.

You cannot use this clause to drop:

- A pseudocolumn, clustered column, or partitioning column. (You can drop nonpartitioning columns from a partitioned table if all the tablespaces where the partitions were created are online and in read-write mode.)
- A column from a nested table, an object table, or a table owned by SYS

*allocate\_extent\_  
clause*

explicitly allocates a new extent for the table, the partition or subpartition, the overflow data segment, the LOB data segment, or the LOB index.

**Restriction:** You cannot allocate an extent for a range- or composite-partitioned table.

**SIZE** specifies the size of the extent in bytes. Use K or M to specify the extent size in kilobytes or megabytes. If you omit this parameter, Oracle determines the size based on the values of the `STORAGE` parameters of the table's overflow data segment or of the LOB index.

**DATAFILE** specifies one of the datafiles in the tablespace of the table, overflow data segment, LOB data tablespace, or LOB index to contain the new extent. If you omit this parameter, Oracle chooses the datafile.

**INSTANCE** makes the new extent available to the freelist group associated with the specified instance. If the instance number exceeds the maximum number of freelist groups, the former is divided by the latter, and the remainder is used to identify the freelist group to be used. An instance is identified by the value of its initialization parameter `INSTANCE_NUMBER`. If you omit this parameter, the space is allocated to the table, but is not drawn from any particular freelist group. Rather, the master freelist is used, and space is allocated as needed. **Use this parameter only if you are using Oracle with the Parallel Server option in parallel mode.**

**See Also:** *Oracle8i Concepts*

Explicitly allocating an extent with this clause does affect the size for the next extent to be allocated as specified by the `NEXT` and `PCTINCREASE` storage parameters.

*deallocate\_  
unused\_clause*

explicitly deallocates unused space at the end of the table, partition or subpartition, overflow data segment, LOB data segment, or LOB index and makes the space available for other segments in the tablespace. You can free only unused space above the high water mark (that is, the point beyond which database blocks have not yet been formatted to receive data).

Oracle credits the amount of the released space to the user quota for the tablespace in which the deallocation occurs.

Oracle deallocates unused space from the end of the object toward the high water mark at the beginning of the object. If an extent is completely contained in the deallocation, then the whole extent is freed for reuse. If an extent is partially contained in the deallocation, then the used part up to the high water mark becomes the extent, and the remaining unused space is freed for reuse.

The exact amount of space freed depends on the values of the `INITIAL`, `MINEXTENTS`, and `NEXT` parameters (as described in "[storage\\_clause](#)" on page 7-605).

`KEEP` specifies the number of bytes above the high water mark that the table, overflow data segment, LOB data segment, or LOB index will have after deallocation.

- If you omit `KEEP` and the high water mark is above the size of `INITIAL` and `MINEXTENTS`, then all unused space above the high water mark is freed. When the high water mark is less than the size of `INITIAL` or `MINEXTENTS`, then all unused space above `MINEXTENTS` is freed.
- If you specify `KEEP`, then the specified amount of space is kept and the remaining space is freed. When the remaining number of extents is less than `MINEXTENTS`, then `MINEXTENTS` is adjusted to the new number of extents. If the initial extent becomes smaller than `INITIAL`, then `INITIAL` is adjusted to the new size.
- In either case, `NEXT` is set to the size of the last extent that was deallocated.

`CACHE` for data that is accessed frequently, specifies that the blocks retrieved for this table are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables.

As a parameter in the *LOB\_storage\_clause*, `CACHE` specifies that Oracle places LOB data values in the buffer cache for faster access.

**Restriction:** You cannot specify `CACHE` for an index-organized table. However, index-organized tables implicitly provide `CACHE` behavior.

`NOCACHE` for data that is not accessed frequently, specifies that the blocks retrieved for this table are placed at the least recently used end of the LRU list in the buffer cache when a full table scan is performed.

As a parameter in the *LOB\_storage\_clause*, `NOCACHE` specifies that the LOB value is either not brought into the buffer cache or brought into the buffer cache and placed at the least recently used end of the LRU list. (The latter is the default behavior.) `NOCACHE` is the default for LOB storage.

**Restriction:** You cannot specify `NOCACHE` for index-organized tables.

**CACHE READS** applies only to LOB storage. It specifies that LOB values are brought into the buffer cache only during read operations, but not during write operations.

- When you add a new LOB column, you can specify the logging attribute with `CACHE READS`, as you can when defining a LOB column at create time.
- When you modify a LOB column from `CACHE` or `NOCACHE` to `CACHE READS`, or from `CACHE READS` to `CACHE` or `NOCACHE`, you can change the logging attribute. If you do not specify the `LOGGING` or `NOLOGGING`, this attribute defaults to the current logging attribute of the LOB column.

For existing LOBs, if you do not specify `CACHE`, `NOCACHE`, or `CACHE READS`, Oracle retains the existing values of the LOB attributes.

**MONITORING** specifies that Oracle can collect modification statistics on *table*. These statistics are estimates of the number of rows affected by DML statements over a particular period of time. They are available for use by the optimizer or for analysis by the user.

**See Also:** *Oracle8i Designing and Tuning for Performance* for more information on using this clause.

**NOMONITORING** specifies that Oracle will not collect modification statistics on *table*.

**Restriction:** You cannot specify `MONITORING` or `NOMONITORING` for a temporary table.

**LOGGING | NOLOGGING** specifies whether subsequent Direct Loader (SQL\*Loader) and direct-load `INSERT` operations against a nonpartitioned table, table partition, all partitions of a partitioned table, or all subpartitions of a partition will be logged (`LOGGING`) or not logged (`NOLOGGING`) in the redo log file.

When used with the *modify\_default\_attributes\_clause*, this clause affects the logging attribute of a partitioned table.

`LOGGING | NOLOGGING` also specifies whether `ALTER TABLE . . . MOVE` and `ALTER TABLE . . . SPLIT` operations will be logged or not logged.

**For a table or table partition**, if you omit `LOGGING | NOLOGGING`, the logging attribute of the table or table partition defaults to the logging attribute of the tablespace in which it resides.

**For LOBs**, if you omit `LOGGING | NOLOGGING`,

- If you specify `CACHE`, then `LOGGING` is used (because you cannot have `CACHE NOLOGGING`).
- If you specify `NOCACHE` or `CACHE READS`, the logging attribute defaults to the logging attribute of the tablespace in which it resides.

`NOLOGGING` does not apply to LOBs that are stored inline with row data. That is, if you specify `NOLOGGING` for LOBs with values less than 4000 bytes and you have not disabled `STORAGE IN ROW`, Oracle ignores the `NOLOGGING` specification and treats the LOB data the same as other table data.

In `NOLOGGING` mode, data is modified with minimal logging (to mark new extents invalid and to record dictionary changes). When applied during media recovery, the extent invalidation records mark a range of blocks as logically corrupt, because the redo data is not logged. Therefore, if you cannot afford to lose this table, it is important to take a backup after the `NOLOGGING` operation.

If the database is run in `ARCHIVELOG` mode, media recovery from a backup taken before the `LOGGING` operation will restore the table. However, media recovery from a backup taken before the `NOLOGGING` operation will not restore the table.

The logging attribute of the base table is independent of that of its indexes.

**See Also:** *Oracle8i Parallel Server Concepts* for more information about the `logging_clause` and parallel DML.

`RENAME TO` renames *table* to *new\_table\_name*.

**Note:** Using this clause will invalidate any dependent materialized views.

**See Also:** "[CREATE MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-318 and *Oracle8i Data Warehousing Guide* for more information on materialized views.

`records_per_block_clause` determines whether Oracle restricts the number of records that can be stored in a block. This clause ensures that any bitmap indexes subsequently created on the table will be as small (compressed) as possible.

**Restrictions:**

- You cannot specify either `MINIMIZE` or `NOMINIMIZE` if a bitmap index has already been defined on *table*. You must first drop the bitmap index.
- You cannot specify this clause for an index-organized table or nested table.

`MINIMIZE` instructs Oracle to calculate the largest number of records in any block in the table, and limit future inserts so that no block can contain more than that number of records.

**Restriction:** You cannot specify `MINIMIZE` for an empty table.

`NOMINIMIZE` disables the `MINIMIZE` feature. This is the default.

`alter_overflow_clause` modifies the definition of an index-organized table. Index-organized tables keep data sorted on the primary key and are therefore best suited for primary-key-based access and manipulation.

**Note:** When you alter an index-organized table, Oracle evaluates the maximum size of each column to estimate the largest possible row. If an overflow segment is needed but you have not specified `OVERFLOW`, Oracle raises an error and does not execute the `ALTER TABLE` statement. This checking function guarantees that subsequent DML operations on the index-organized table will not fail because an overflow segment is lacking.

<code>PCTTHRESHOLD</code> <i>integer</i>	<p>specifies the percentage of space reserved in the index block for an index-organized table row. Any portion of the row that exceeds the specified threshold is stored in the overflow area. <code>PCTTHRESHOLD</code> must be a value from 1 to 50.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot reduce the value of <code>PCTTHRESHOLD</code> so much that the primary key will not fit.</li><li>■ You cannot specify <code>PCTTHRESHOLD</code> for individual partitions of an index-organized table.</li></ul>
<code>INCLUDING</code> <i>column_name</i>	<p>specifies the column at which to divide an index-organized table row into index and overflow portions. All non-primary-key columns that follow <i>column_name</i> are stored in the overflow data segment. The <i>column_name</i> is either the name of the last primary key column or any subsequent non-primary-key column.</p> <p>If you use the <i>drop_column_clause</i> to drop (or mark unused) a column defined as an <code>INCLUDING</code> column, the column stored immediately before this column will become the new <code>INCLUDING</code> column.</p>
<i>overflow_clause</i>	<p>specifies the overflow data segment physical storage and logging attributes to be modified for the index-organized table. Parameters specified in this clause are applicable only to the overflow data segment.</p> <p><b>Restriction:</b> You cannot specify <code>OVERFLOW</code> for a partition of a partitioned index-organized table unless the table already has an overflow segment.</p> <p><b>See Also:</b> "<a href="#">CREATE TABLE</a>" on page 4-381.</p>

---

<i>add_overflow_clause</i>	<p>adds an overflow data segment to the specified index-organized table.</p> <p>For a partitioned index-organized table:</p> <ul style="list-style-type: none"> <li>■ If you do not specify <code>PARTITION</code>, Oracle automatically allocates an overflow segment for each partition. The physical attributes of these segments are inherited from the table level.</li> <li>■ If you wish to specify separate physical attributes for one or more partitions, you must specify such attributes for <b>every</b> partition in the table. You do not specify the name of the partitions, but you must specify their attributes in the order in which they were created.</li> </ul> <p>You can find the order of the partitions by querying the <code>PARTITION_NAME</code> and <code>PARTITION_POSITION</code> columns of the <code>USER_IND_PARTITIONS</code> view.</p> <p>If you do not specify <code>TABLESPACE</code> for a particular partition, Oracle uses the tablespace specified for the table. If you do not specify <code>TABLESPACE</code> at the table level, Oracle uses the tablespace of the partition's primary key index segment.</p>
<i>partitioning_clauses</i>	<p>The following clauses apply only to partitioned tables. You cannot combine partition operations with other partition operations or with operations on the base table in one <code>ALTER TABLE</code> statement.</p> <hr/> <p><b>Note:</b> If you drop, exchange, truncate, move, modify, or split a partition on a table that is a master table for one or more materialized views, existing bulk load information about the table will be deleted. Therefore, be sure to refresh all dependent materialized views before performing any of these operations.</p>
<i>modify_default_attributes_clause</i>	<p>specifies new default values for the attributes of <i>table</i>. Partitions and LOB partitions you create subsequently will inherit these values unless you override them explicitly when creating the partition or LOB partition. Existing partitions and LOB partitions are not affected by this clause.</p> <p>Only attributes named in the statement are affected, and the default values specified are overridden by any attributes specified at the individual partition level.</p> <p><code>FOR PARTITION</code> applies only to composite-partitioned tables. This clause specifies new default values for the attributes of <i>partition</i>. Subpartitions and LOB subpartitions of <i>partition</i> that you create subsequently will inherit these values, unless you override them explicitly when creating the subpartition or LOB subpartition. Existing subpartitions are not affected by this clause.</p>

**Restrictions:**

- The PCTTHRESHOLD, COMPRESS, *physical\_attributes\_clause*, and *overflow\_clause* are valid only for partitioned index-organized tables.
- You cannot specify the PCTUSED parameter for the index segment of an index-organized table.
- You can specify COMPRESS only if compression is already specified at the table level.

*modify\_partition\_clause*

modifies the real physical attributes of the *partition* table partition. Optionally modifies the storage attributes of one or more LOB items for the partition. You can specify new values for any of the following physical attributes for the partition: the logging attribute; PCTFREE, PCTUSED, INITRANS, or MAXTRANS parameter; or storage parameters.

If *table* is composite-partitioned:

- If you specify the *allocate\_extent\_clause*, Oracle will allocate an extent for each subpartition of *partition*.
- If you specify *deallocate\_unused\_clause*, Oracle will deallocate unused storage from each subpartition of *partition*.
- Any other attributes changed in this clause will be changed in subpartitions of *partition* as well, overriding existing values. To avoid changing the attributes of existing subpartitions, use the FOR PARTITION clause of the *modify\_default\_attributes\_clause*.

**Restriction:** If *table* is hash partitioned, you can specify only the *allocate\_extent* and *deallocate\_unused* clauses. All other attributes of the partition are inherited from the table-level defaults except TABLESPACE, which stays the same as it was at create time.

*add\_subpartition\_clause* adds a hash subpartition to *partition*. Oracle populates the new subpartition with rows rehashed from the other subpartition(s) of *partition* as determined by the hash function.

Oracle marks UNUSABLE, and you must rebuild, the local index subpartitions corresponding to the added and to the rehashed subpartitions.

If you do not specify *subpartition*, Oracle assigns a name in the form SYS\_SUBPnnn

If you do not specify TABLESPACE, the new subpartition will reside in the default tablespace of *partition*.

COALESCE  
SUBPARTITION

specifies that Oracle should select a hash subpartition, distribute its contents into one or more remaining subpartitions (determined by the hash function), and then drop the selected subpartition.

Local index subpartitions corresponding to the selected subpartition are also dropped. Oracle marks UNUSABLE, and you must rebuild, the index subpartitions corresponding to one or more absorbing subpartitions.



---

UNUSABLE LOCAL INDEXES clause	<p>The next two clauses modify the attributes of local <b>index partitions</b> corresponding to <i>partition</i>.</p> <p>UNUSABLE LOCAL INDEXES marks UNUSABLE all the local index partitions associated with <i>partition</i>.</p> <p>REBUILD UNUSABLE LOCAL INDEXES rebuilds the unusable local index partitions associated with <i>partition</i>.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify this clause with any other clauses of the <i>modify_partition_clause</i>.</li> <li>■ You cannot specify this clause for partitions that are subpartitioned.</li> </ul>
<i>modify_subpartition_clause</i>	<p>lets you allocate or deallocate storage for an individual subpartition of <i>table</i>.</p> <p><b>Restriction:</b> The only <i>modify_LOB_storage_parameters</i> you can specify for subpartition are the <i>allocate_extent_clause</i> and <i>deallocate_unused_clause</i>.</p> <p>UNUSABLE LOCAL INDEXES marks UNUSABLE all the local index subpartitions associated with <i>subpartition</i>.</p> <p>REBUILD UNUSABLE LOCAL INDEXES rebuilds the unusable local index subpartitions associated with <i>subpartition</i>.</p>
<i>rename_partition/subpartition_clause</i>	<p>renames a table partition or subpartition <i>current_name</i> to <i>new_name</i>. For both partitions and subpartitions, <i>new_name</i> must be different from all existing partitions and subpartitions of the same table.</p>
<i>move_partition_clause</i>	<p>moves table partition <i>partition</i> to another segment. You can move partition data to another tablespace, recluster data to reduce fragmentation, or change create-time physical attributes.</p> <p>If the table contains LOB columns, you can use the <i>LOB_storage_clause</i> to move the LOB data and LOB index segments associated with this partition. Only the LOBs named are affected. If you do not specify the <i>LOB_storage_clause</i> for a particular LOB column, its LOB data and LOB index segments are not moved.</p> <p>If <i>partition</i> is not empty, MOVE PARTITION marks UNUSABLE all corresponding local index partitions and all global nonpartitioned indexes, and all the partitions of global partitioned indexes.</p> <p>When you move a LOB data segment, Oracle drops the old data segment and corresponding index segment and creates new segments even if you do not specify a new tablespace.</p>

The move operation obtains its parallel attribute from the *parallel\_clause*, if specified. If not specified, the default parallel attributes of the table, if any, are used. If neither is specified, Oracle performs the move without using parallelism.

The *parallel\_clause* on `MOVE PARTITION` does not change the default parallel attributes of *table*.

---

**Note:** For index-organized tables, Oracle uses the address of the primary key, as well as its value, to construct logical rowids. The logical rowids are stored in the secondary index of the table. If you move a partition of an index-organized table, the address portion of the rowids will change, which can hamper performance. To ensure optimal performance, rebuild the secondary index(es) on the moved partition to update the rowids.

**See Also:** *Oracle8i Concepts* for more information on logical rowids.

---

**Restrictions:**

- If *partition* is a hash partition, the only attribute you can specify in this clause is `TABLESPACE`.
- You cannot move a partition of a composite-partitioned table. You must move each subpartition separately with the *move\_subpartition\_clause*.
- You cannot specify this clause for a partition containing subpartitions. However, you can move subpartitions using the *move\_subpartition\_clause*.

*move\_subpartition\_clause*

moves the table subpartition *subpartition* to another segment. If you do not specify `TABLESPACE`, the subpartition will remain in the same tablespace.

Unless the subpartition is empty, Oracle marks `UNUSABLE` all local index subpartitions corresponding to the subpartition being moved, as well as global nonpartitioned indexes and partitions of global indexes.

If the table contains LOB columns, you can use the *LOB\_storage\_clause* to move the LOB data and LOB index segments associated with this subpartition. Only the LOBs named are affected. If you do not specify the *LOB\_storage\_clause* for a particular LOB column, its LOB data and LOB index segments are not moved.

When you move a LOB data segment, Oracle drops the old data segment and corresponding index segment and creates new segments even if you do not specify a new tablespace.

*add\_range\_partition\_clause*

adds a new range partition *partition* to the "high" end of a partitioned table (after the last existing partition). You can specify any create-time physical attributes for the new partition. If the table contains LOB columns, you can also specify partition-level attributes for one or more LOB items.

You can specify up to 64K-1 partitions. For a discussion of factors that might impose practical limits less than this number, refer to *Oracle8i Administrator's Guide*.

**Restrictions:**

- If the first element of the partition bound of the high partition is `MAXVALUE`, you cannot add a partition to the table. Instead, use the *split\_partition\_clause* to add a partition at the beginning or the middle of the table.
- The *compression\_clause*, *physical\_attributes\_clause*, and `OVERFLOW` are valid only for a partitioned index-organized table.
- You cannot specify the `PCTUSED` parameter for the index segment of an index-organized table.
- You can specify `OVERFLOW` only if the partitioned table already has an overflow segment.
- You can specify compression only if compression is enabled at the table level.

`VALUES LESS THAN (value_list)` specifies the upper bound for the new partition. The *value\_list* is a comma-separated, ordered list of literal values corresponding to *column\_list*. The *value\_list* must collate greater than the partition bound for the highest existing partition in the table.

*partition\_level\_subpartitioning* is permitted only for a composite-partitioned table. This clause lets you specify particular hash subpartitions for *partition*. You specify composite partitioning in one of two ways:

- You can specify individual subpartitions by name, and optionally the tablespace where each should be stored, or
- You can specify the number of subpartitions (and optionally one or more tablespaces where they are to be stored). In this case, Oracle assigns partition names of the form `SYS_SUBPnnn`. The number of tablespaces does not have to equal the number of subpartitions. If the number of subpartitions is greater than the number of tablespaces, Oracle cycles through the names of the tablespaces.

The subpartitions inherit all their attributes from any attributes specified for *new\_partition*, except for `TABLESPACE`, which you can specify at the subpartition level. Any attributes not specified at the subpartition or partition level are inherited from table-level defaults.

This clause overrides any subpartitioning specified at the table level.

If you do not specify this clause but you specified default subpartitioning at the table level, *new\_partition\_name* will inherit the table-level default subpartitioning (see ["CREATE TABLE"](#) on page 4-381).

---

<i>add_hash_partition_clause</i>	<p>adds a new hash partition to the "high" end of a partitioned table. Oracle will populate the new partition with rows rehashed from other partitions of <i>table</i> as determined by the hash function.</p> <p>You can specify a name for the partition, and optionally a tablespace where it should be stored. If you do not specify <i>new_partition_name</i>, Oracle assigns a partition name of the form <i>SYS_Pnnn</i>. If you do not specify <i>TABLESPACE</i>, the new partition is stored in the table's default tablespace. Other attributes are always inherited from table-level defaults.</p> <p><b>See Also:</b> "<a href="#">CREATE TABLE</a>" on page 4-381 and <i>Oracle8i Concepts</i> for more information on hash partitioning.</p>
<i>parallel_clause</i>	<p>lets you specify whether to parallelize the creation of the new partition.</p>
<i>coalesce_partition_clause</i>	<p>applies only to hash-partitioned tables. This clause specifies that Oracle should select a hash partition, distribute its contents into one or more remaining partitions (determined by the hash function), and then drop the selected partition. Local index partitions corresponding to the selected partition are also dropped. Oracle marks <i>UNUSABLE</i>, and you must rebuild, the local index partitions corresponding to one or more absorbing partitions.</p>
<i>drop_partition_clause</i>	<p>applies only to tables partitioned using the range or composite method. This clause removes partition <i>partition</i>, and the data in that partition, from a partitioned table. If you want to drop a partition but keep its data in the table, you must merge the partition into one of the adjacent partitions. See the <a href="#">merge_partitions_clause</a> of this statement.</p> <p>If the table has LOB columns, the LOB data and LOB index partitions (and their subpartitions, if any) corresponding to <i>partition</i> are also dropped.</p> <ul style="list-style-type: none"><li>■ Oracle drops local index partitions and subpartitions corresponding to <i>partition</i>, even if they are marked <i>UNUSABLE</i>.</li><li>■ Oracle marks <i>UNUSABLE</i> all global nonpartitioned indexes defined on the table and all partitions of global partitioned indexes, unless the partition being dropped or all of its subpartitions are empty.</li><li>■ If you drop a partition and later insert a row that would have belonged to the dropped partition, Oracle stores the row in the next higher partition. However, if that partition is the highest partition, the insert will fail because the range of values represented by the dropped partition is no longer valid for the table.</li></ul> <p><b>Restriction:</b> If <i>table</i> contains only one partition, you cannot drop the partition. You must drop the table.</p>
<i>truncate_partition_clause</i>	<p><i>PARTITION</i> removes all rows from <i>partition</i> or, if the table is composite-partitioned, all rows from <i>partition</i>'s subpartitions. <i>SUBPARTITION</i> removes all rows from <i>subpartition</i>.</p>
<i>truncate_subpartition_clause</i>	<p>If the table contains any LOB columns, the LOB data and LOB index segments for this partition are also truncated. If the table is composite-partitioned, the LOB data and LOB index segments for this partition's subpartitions are truncated.</p>

If the partition or subpartition to be truncated contains data, you must first disable any referential integrity constraints on the table. Alternatively, you can delete the rows and then truncate the partition.

For each partition or subpartition truncated, Oracle also truncates corresponding local index partitions and subpartitions. If those index partitions or subpartitions are marked UNUSABLE, Oracle truncates them and resets the UNUSABLE marker to VALID. In addition, if the truncated partition or subpartition, or any of the subpartitions of the truncated partition are not empty, Oracle marks as UNUSABLE all global nonpartitioned indexes and partitions of global indexes defined on the table.

DROP STORAGE deallocates space from the deleted rows and makes it available for use by other schema objects in the tablespace.

REUSE STORAGE keeps space from the deleted rows allocated to the partition or subpartition. The space is subsequently available only for inserts and updates to the same partition or subpartition.

*split\_partition\_*  
*clause*

from an original partition *partition\_name\_old*, creates two new partitions, each with a new segment and new physical attributes, and new initial extents. The segment associated with *partition\_name\_old* is discarded.

**Restriction:** You cannot specify this clause for a hash-partitioned table.

AT (*value\_*  
*list*) specifies the new noninclusive upper bound for *split\_partition\_1*. The *value\_list* must compare less than the original partition bound for *partition\_name\_old* and greater than the partition bound for the next lowest partition (if there is one).

INTO describes the two partitions resulting from the split.

*partition\_*  
*description*,  
*partition\_*  
*description* specifies optional names and physical attributes of the two partitions resulting from the split. If you do not specify new partition names, Oracle assigns names of the form SYS\_P*n*. Any attributes you do not specify are inherited from *partition\_name\_old*.

**Restriction:**

- You can specify the *compression\_clause*, *physical\_attributes\_clause*, and OVERFLOW only for a partitioned index-organized table.
- You cannot specify the PCTUSED parameter for the index segment of an index-organized table.

*parallel\_clause* parallelizes the split operation, but does not change the default parallel attributes of the table.

If you specify subpartitioning for the new partitions, you can specify only TABLESPACE for the subpartitions. All other attributes will be inherited from the containing new partition.

If *partition\_name\_old* is subpartitioned, and you do not specify any subpartitioning for the new partitions, the new partitions will inherit the number and tablespaces of the subpartitions in *partition\_name\_old*.

Oracle also splits corresponding local index partitions, even if they are marked `UNUSABLE`. The resulting local index partitions inherit all their partition-level default attributes from the local index partition being split.

If `partition_name_old` was not empty, Oracle marks `UNUSABLE` all global nonpartitioned indexes and all partitions of global indexes on the table. (This action on global indexes does not apply to index-organized tables.) In addition, if any partitions or subpartitions resulting from the split are not empty, Oracle marks as `UNUSABLE` all corresponding local index partitions and subpartitions.

If `table` contains LOB columns, you can use the `LOB_storage_clause` to specify separate LOB storage attributes for the LOB data segments resulting from the split. Oracle drops the LOB data and LOB index segments of `partition_name_old` and creates new segments for each LOB column, for each partition, even if you do not specify a new tablespace.

`merge_partitions_clause` merges the contents of two adjacent partitions of `table` into one new partition, and then drops the original two partitions.

The new partition inherits the partition-bound of the higher of the two original partitions.

Any attributes not specified in the `segment_attributes_clause` are inherited from table-level defaults.

If you do not specify `new_partition_name`, Oracle assigns a name of the form `SYS_Pnnn`. If the new partition has subpartitions, Oracle assigns subpartition names of the form `SYS_SUBPnnn`.

If either or both of the original partitions was not empty, Oracle marks `UNUSABLE` all global nonpartitioned global indexes and all partitions of global indexes on the table. In addition, if the partition or any of its subpartitions resulting from the merge is not empty, Oracle marks `UNUSABLE` all corresponding local index partitions and subpartitions.

**Restriction:** You cannot specify this clause for an index-organized table or for a table partitioned using the hash method.

`partition_level_partitioning` specifies hash subpartitioning attributes for the new partition. Any attributes not specified in this clause are inherited from table-level defaults.

If you do not specify this clause, the new merged partition inherits subpartitioning attributes from table-level defaults.

`parallel_clause` specifies that the merging operation is to be parallelized.

`exchange_partition_clause` lets you exchange the data and index segments of

- a hash or range partition (or subpartition) with a nonpartitioned table
- a hash-partitioned table with a range partition of a composite-partitioned table

`exchange_subpartition_clause`

The default behavior is `EXCLUDING INDEXES WITH VALIDATION`. You must have `ALTER TABLE` privileges on **both** tables to perform this operation.

This clause facilitates high-speed data loading when used with transportable tablespaces. For information on this topic, see *Oracle8i Administrator's Guide*.

If *table* contains LOB columns, for each LOB column Oracle exchanges LOB data and LOB index partition or subpartition segments with corresponding LOB data and LOB index segments of *table*.

All statistics of the table and partition are exchanged, including table, column, index statistics, and histograms. The aggregate statistics of the table receiving the new partition are recalculated.

The logging attribute of the table and partition is also exchanged.

**Restriction:** Both tables involved in the exchange must have the same primary key, and no validated foreign keys can be referencing either of the tables unless the referenced table is empty.

WITH TABLE <i>table</i>	specifies the table with which the partition will be exchanged.
INCLUDING INDEXES	specifies that the local index partitions or subpartitions should be exchanged with the corresponding table index (for a nonpartitioned table) or local indexes (for a hash-partitioned table).
EXCLUDING INDEXES	specifies that all index partitions or subpartitions corresponding to the partition and all the regular indexes and index partitions on the exchanged table are marked UNUSABLE.
WITH VALIDATION	specifies that if any rows in the exchanged table do not map into partitions or subpartitions being exchanged, Oracle should return an error.
WITHOUT VALIDATION	specifies that the proper mapping of rows in the exchanged table is not checked.
EXCEPTIONS INTO	specifies a table into which Oracle places the rowids of all rows violating the constraint. If you omit <i>schema</i> , Oracle assumes the exceptions table is in your own schema. If you omit this clause altogether, Oracle assumes that the table is named EXCEPTIONS. The exceptions table must be on your local database.

You can create the EXCEPTIONS table using one of these scripts:

- `UTLEXCEPT.SQL` uses physical rowids. Therefore it can accommodate rows from conventional tables but not from index-organized tables. (See the Note that follows.)
- `UTLEXPT1.SQL` uses universal rowids, so it can accommodate rows from both conventional and index-organized tables.

If you create your own exceptions table, it must follow the format prescribed by one of these two scripts. See *Oracle8i Migration* for compatibility issues related to the use of these scripts.

**Note:** If you are collecting exceptions from index-organized tables based on primary keys (rather than universal rowids), you must create a separate exceptions table for each index-organized table to accommodate its primary-key storage. You create multiple exceptions tables with different names by modifying and resubmitting the script.

For information on the SQL scripts, see the `DBMS_IOT` package in *Oracle8i Supplied PL/SQL Packages Reference*. For information on eliminating migrated and chained rows, see *Oracle8i Designing and Tuning for Performance*.

**Restrictions on EXCEPTIONS INTO clause:**

- This clause is not valid with subpartitions.
- The partitioned table must have been defined with a `UNIQUE` constraint, and that constraint must be in `DISABLE VALIDATE` state.

If these conditions are not true, Oracle ignores this clause.

**See Also:** The "[constraint\\_clause](#)" on page 7-233 for more information on constraint checking.

**Restrictions on exchanging partitions:**

When exchanging between a hash-partitioned table and the range partition of a composite-partitioned table, the following restrictions apply:

- The partitioning key of the hash-partitioned table must be identical to the subpartitioning key of the composite-partitioned table.
- The number of partitions in the hash-partitioned table must be identical to the number of subpartitions in the range partition of the composite-partitioned table.
- Oracle marks `UNUSABLE` all global indexes on both tables.

For partitioned index-organized tables, the following additional restrictions apply:

- The source and target table/partition must have their primary key set on the same columns, in the same order.
- If compression is enabled, it must be enabled for both the source and the target, and with the same prefix length.
- An index-organized table partition cannot be exchanged with a regular table or vice versa.
- Both the source and target must have overflow segments, or neither can have overflow segments.

*row\_movement\_clause* determines whether a row can be moved to a different partition or subpartition because of a change to one or more of its key values.

**Restriction:** You can specify this clause only for a partitioned table.

`ENABLE` allows Oracle to move a row to a different partition or subpartition as the result of an update to the partitioning or subpartitioning key.

**Restriction:** You cannot specify this clause if a domain index has been built on any column of the table.

---

**WARNING: Moving a row in the course of an UPDATE operation changes that row's rowid.**

---

`DISABLE` returns an error if an update to a partitioning or subpartitioning key would result in a row moving to a different partition or subpartition. This is the default.

*parallel\_clause* changes the default degree of parallelism for queries and DML on the table. For additional information, see the [Notes](#) to the *parallel\_clause* of "[CREATE TABLE](#)" on page 4-381.



---

NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the PARALLEL_THREADS_PER_CPU initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution processes. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .
<p><b>Restriction:</b> If <i>table</i> contains any columns of LOB or user-defined object type, subsequent INSERT, UPDATE, and DELETE operations on <i>table</i> are executed serially without notification. Subsequent queries, however, will be executed in parallel.</p>	
<p><b>Note:</b> If you specify the <i>parallel_clause</i> in conjunction with the <i>move_table_clause</i>, the parallelism applies only to the move, not to subsequent DML and query operations on the table.</p>	
<i>enable_disable_clause</i>	lets you specify whether Oracle should apply an integrity constraint. For a complete description of this clause, including notes and restrictions that relate to this statement, see the <a href="#">enable_disable_clause</a> of "CREATE TABLE" on page 4-381.
ENABLE TABLE LOCK	enables DML and DDL locks on a table in a parallel server environment. <b>See Also:</b> <i>Oracle8i Parallel Server Concepts</i> .
<p><b>Note:</b> DML table locks are not acquired on temporary tables.</p>	
DISABLE TABLE LOCK	disables DML and DDL locks on a table to improve performance in a parallel server environment. <b>See Also:</b> <i>Oracle8i Parallel Server Concepts</i> .
ENABLE ALL TRIGGERS	enables all triggers associated with the table. Oracle fires the triggers whenever their triggering condition is satisfied. See " <a href="#">CREATE TRIGGER</a> " on page 7-426.  To enable a single trigger, use the <i>enable_clause</i> of ALTER TRIGGER. See " <a href="#">ALTER TRIGGER</a> " on page 7-183.
DISABLE ALL TRIGGERS	disables all triggers associated with the table. Oracle will not fire a disabled trigger even if the triggering condition is satisfied.

---

## Examples

**Nested Table Example** The following statement modifies the storage characteristics of a nested table column PROJECTS in table EMP so that when queried it returns actual values instead of locators:

```
ALTER TABLE emp MODIFY NESTED TABLE projects RETURN AS VALUE;
```

**PARALLEL Example** The following statement specifies parallel processing for queries to the EMP table:

```
ALTER TABLE emp
  PARALLEL;
```

**ENABLE VALIDATE Example** The following statement places in ENABLE VALIDATE state an integrity constraint named FK\_DEPTNO in the EMP table:

```
ALTER TABLE emp
  ENABLE VALIDATE CONSTRAINT fk_deptno
  EXCEPTIONS INTO except_table;
```

Each row of the EMP table must satisfy the constraint for Oracle to enable the constraint. If any row violates the constraint, the constraint remains disabled. Oracle lists any exceptions in the table EXCEPT\_TABLE. You can also identify the exceptions in the EMP table with the following statement:

```
SELECT emp.*
  FROM emp e, except_table ex
 WHERE e.row_id = ex.row_id
       AND ex.table_name = 'EMP'
       AND ex.constraint = 'FK_DEPTNO';
```

**ENABLE NOVALIDATE Example** The following statement tries to place in ENABLE NOVALIDATE state two constraints on the EMP table:

```
ALTER TABLE emp
  ENABLE NOVALIDATE UNIQUE (ename)
  ENABLE NOVALIDATE CONSTRAINT nn_ename;
```

This statement has two ENABLE clauses:

- The first places a unique constraint on the ENAME column in ENABLE NOVALIDATE state.
- The second places the constraint named NN\_ENAME in ENABLE NOVALIDATE state.

In this case, Oracle enables the constraints only if both are satisfied by each row in the table. If any row violates either constraint, Oracle returns an error and both constraints remain disabled.

**DISABLE Example** Consider a referential integrity constraint involving a foreign key on the combination of the AREACO and PHONENO columns of the PHONE\_CALLS table. The foreign key references a unique key on the combination of the AREACO

and PHONENO columns of the CUSTOMERS table. The following statement disables the unique key on the combination of the AREACO and PHONENO columns of the CUSTOMERS table:

```
ALTER TABLE customers
    DISABLE UNIQUE (areaco, phoneno) CASCADE;
```

The unique key in the CUSTOMERS table is referenced by the foreign key in the PHONE\_CALLS table, so you must use the CASCADE clause to disable the unique key. This clause disables the foreign key as well.

**CHECK Example** The following statement defines and disables a CHECK constraint on the EMP table:

```
ALTER TABLE emp
    ADD (CONSTRAINT check_comp CHECK (sal + comm <= 5000) )
    DISABLE CONSTRAINT check_comp;
```

The constraint CHECK\_COMP ensures that no employee's total compensation exceeds \$5000. The constraint is disabled, so you can increase an employee's compensation above this limit.

**Triggers Example** The following statement enables all triggers associated with the EMP table:

```
ALTER TABLE emp
    ENABLE ALL TRIGGERS;
```

**DEALLOCATE UNUSED Example** The following statement frees all unused space for reuse in table EMP, where the high water mark is above MINEXTENTS:

```
ALTER TABLE emp
    DEALLOCATE UNUSED;
```

**DROP COLUMN Example** This statement illustrates the *drop\_column\_clause* with CASCADE CONSTRAINTS. Assume table T1 is created as follows:

```
CREATE TABLE t1 (
    pk NUMBER PRIMARY KEY,
    fk NUMBER,
    c1 NUMBER,
    c2 NUMBER,
    CONSTRAINT ri FOREIGN KEY (fk) REFERENCES t1,
    CONSTRAINT ck1 CHECK (pk > 0 and c1 > 0),
    CONSTRAINT ck2 CHECK (c2 > 0)
```

```
);
```

An error will be returned for the following statements:

```
ALTER TABLE t1 DROP (pk); -- pk is a parent key
ALTER TABLE t1 DROP (c1); -- c1 is referenced by multicolumn
                           constraint ck1
```

Submitting the following statement drops column PK, the primary key constraint, the foreign key constraint, RI, and the check constraint, CK1:

```
ALTER TABLE t1 DROP (pk) CASCADE CONSTRAINTS;
```

If all columns referenced by the constraints defined on the dropped columns are also dropped, then `CASCADE CONSTRAINTS` is not required. For example, assuming that no other referential constraints from other tables refer to column PK, then it is valid to submit the following statement without the `CASCADE CONSTRAINTS` clause:

```
ALTER TABLE t1 DROP (pk, fk, c1);
```

**Index-Organized Table Examples** This statement modifies the `INITRANS` parameter for the index segment of index-organized table `DOCINDEX`:

```
ALTER TABLE docindex INITRANS 4;
```

The following statement adds an overflow data segment to index-organized table `DOCINDEX`:

```
ALTER TABLE docindex ADD OVERFLOW;
```

This statement modifies the `INITRANS` parameter for the overflow data segment of index-organized table `DOCINDEX`:

```
ALTER TABLE docindex OVERFLOW INITRANS 4;
```

**ADD PARTITION Example** The following statement adds a partition P3 and specifies storage characteristics for three of the table's LOB columns (B, C, and D):

```
ALTER TABLE pt ADD PARTITION p3 VALUES LESS THAN (30)
  LOB (b, d) STORE AS (TABLESPACE tsz)
  LOB (c) STORE AS mylobseg;
```

The LOB data and LOB index segments for columns B and D in partition P3 will reside in tablespace TSZ. The remaining attributes for these LOB columns will be inherited first from the table-level defaults, and then from the tablespace defaults.

The LOB data segments for column C will reside in the MYLOBSEG segment, and will inherit all other attributes from the table-level defaults and then from the tablespace defaults.

**SPLIT PARTITION Example** The following statement splits partition P3 into partitions P3\_1 and P3\_2:

```
ALTER TABLE pt SPLIT PARTITION p3 AT (25)
  INTO (PARTITION p3_1 TABLESPACE ts4
        LOB (b,d) STORE AS (TABLESPACE tsz),
        PARTITION p3_2 (TABLESPACE ts5)
        LOB (c) STORE AS (TABLESPACE ts5));
```

In partition P3\_1, Oracle creates the LOB segments for columns B and D in tablespace TSZ. In partition P3\_2, Oracle creates the LOB segments for column C in tablespace TS5. The LOB segments for columns B and D in partition P3\_2 and those for column C in partition P3\_1 remain in original tablespace for the original partition P3. However, Oracle creates new segments for all the LOB data and LOB index segments, even though they are not moved to a new tablespace.

**User-Defined Object Identifier Example** The following statements create an object type, a corresponding object table with a primary-key-based object identifier, and a table having a user-defined REF column:

```
CREATE TYPE emp_t AS OBJECT (empno NUMBER, address CHAR(30));

CREATE TABLE emp OF emp_t (
  empno PRIMARY KEY)
  OBJECT IDENTIFIER IS PRIMARY KEY;

CREATE TABLE dept (dno NUMBER, mgr_ref REF emp_t SCOPE is emp);
```

The next statements add a constraint and a user-defined REF column, both of which reference table EMP:

```
ALTER TABLE dept ADD CONSTRAINT mgr_cons FOREIGN KEY (mgr_ref)
  REFERENCES emp;
ALTER TABLE dept ADD sr_mgr REF emp_t REFERENCES emp;
```

**Add Column Example** The following statement adds a column named THRIFTPLAN of datatype NUMBER with a maximum of seven digits and two decimal places and a column named LOANCODE of datatype CHAR with a size of one and a NOT NULL integrity constraint:

```
ALTER TABLE emp
```

```
ADD (thriftplan NUMBER(7,2),
     loancode CHAR(1) NOT NULL);
```

**Modify Column Examples** The following statement increases the size of the THRIFTPLAN column to nine digits:

```
ALTER TABLE emp
  MODIFY (thriftplan NUMBER(9,2));
```

Because the MODIFY clause contains only one column definition, the parentheses around the definition are optional.

The following statement changes the values of the PCTFREE and PCTUSED parameters for the EMP table to 30 and 60, respectively:

```
ALTER TABLE emp
  PCTFREE 30
  PCTUSED 60;
```

**ALLOCATE EXTENT Example** The following statement allocates an extent of 5 kilobytes for the EMP table and makes it available to instance 4:

```
ALTER TABLE emp
  ALLOCATE EXTENT (SIZE 5K INSTANCE 4);
```

Because this statement omits the DATAFILE parameter, Oracle allocates the extent in one of the datafiles belonging to the tablespace containing the table.

**DEFAULT Examples** This statement modifies the BAL column of the ACCOUNTS table so that it has a default value of 0:

```
ALTER TABLE accounts
  MODIFY (bal DEFAULT 0);
```

If you subsequently add a new row to the ACCOUNTS table and do not specify a value for the BAL column, the value of the BAL column is automatically 0:

```
INSERT INTO accounts(accno, accname)
  VALUES (accseq.nextval, 'LEWIS');
```

```
SELECT *
  FROM accounts
  WHERE accname = 'LEWIS';
```

```
ACCNO  ACCNAME  BAL
-----  -
```

```
815234 LEWIS      0
```

To discontinue previously specified default values, so that they are no longer automatically inserted into newly added rows, replace the values with nulls, as shown in this statement:

```
ALTER TABLE accounts
    MODIFY (bal DEFAULT NULL);
```

The `MODIFY` clause need only specify the column name and the modified part of the definition, rather than the entire column definition. This statement has no effect on any existing values in existing rows.

**Drop Constraint Examples** The following statement drops the primary key of the `DEPT` table:

```
ALTER TABLE dept
    DROP PRIMARY KEY CASCADE;
```

If you know that the name of the `PRIMARY KEY` constraint is `PK_DEPT`, you could also drop it with the following statement:

```
ALTER TABLE dept
    DROP CONSTRAINT pk_dept CASCADE;
```

The `CASCADE` clause drops any foreign keys that reference the primary key.

The following statement drops the unique key on the `DNAME` column of the `DEPT` table:

```
ALTER TABLE dept
    DROP UNIQUE (dname);
```

The `DROP` clause in this statement omits the `CASCADE` clause. Because of this omission, Oracle does not drop the unique key if any foreign key references it.

**LOB Examples** The following statement adds `CLOB` column `RESUME` to the `EMPLOYEE` table and specifies `LOB` storage characteristics for the new column:

```
ALTER TABLE employee ADD (resume CLOB)
    LOB (resume) STORE AS resume_seg (TABLESPACE resume_ts);
```

To modify the `LOB` column `RESUME` to use caching, enter the following statement:

```
ALTER TABLE employee MODIFY LOB (resume) (CACHE);
```

**Nested Table Examples** The following statement adds the nested table column SKILLS to the EMPLOYEE table:

```
ALTER TABLE employee ADD (skills skill_table_type)
    NESTED TABLE skills STORE AS nested_skill_table;
```

You can also modify a nested table's storage characteristics. Use the name of the storage table specified in the *nested\_table\_storage\_clause* to make the modification. You *cannot* query or perform DML statements on the storage table. Use the storage table only to modify the nested table column storage characteristics.

The following statement creates table VETSERVICE with nested table column CLIENT and storage table CLIENT\_TAB. Nested table VETSERVICE is modified to specify constraints:

```
CREATE TYPE pet_table AS OBJECT
    (pet_name VARCHAR2(10), pet_dob DATE);

CREATE TABLE vetservice (vet_name VARCHAR2(30),
    client pet_table)
    NESTED TABLE client STORE AS client_tab;

ALTER TABLE client_tab ADD UNIQUE (ssn);
```

The following statement adds a UNIQUE constraint to nested table NESTED\_SKILL\_TABLE:

```
ALTER TABLE nested_skill_table ADD UNIQUE (a);
```

The following statement alters the storage table for a nested table of REF values to specify that the REF is scoped:

```
CREATE TYPE emp_t AS OBJECT (eno number, ename char(31));
CREATE TYPE emps_t AS TABLE OF REF emp_t;
CREATE TABLE emptab OF emp_t;
CREATE TABLE dept (dno NUMBER, employees emps_t)
    NESTED TABLE employees STORE AS deptemps;
ALTER TABLE deptemps ADD (SCOPE FOR (column_value) IS emptab);
```

Similarly, to specify storing the REF with rowid:

```
ALTER TABLE deptemps ADD (REF(column_value) WITH ROWID);
```

In order to execute these ALTER TABLE statements successfully, the storage table DEPTEMPS must be empty. Also, because the nested table is defined as a table of



scalars (REFs), Oracle implicitly provides the column name `COLUMN_VALUE` for the storage table.

**See Also:**

- ["CREATE TABLE"](#) on page 4-381 for more information about nested table storage.
- *Oracle8i Application Developer's Guide - Fundamentals* for more information about nested tables.

**REF Examples** In the following statement an object type `DEPT_T` has been previously defined. Now, create table `EMP` as follows:

```
CREATE TABLE emp
  (name VARCHAR(100),
   salary NUMBER,
   dept REF dept_t);
```

An object table `DEPARTMENTS` is created as:

```
CREATE TABLE departments OF dept_t;
```

The `DEPT` column can store references to objects of `DEPT_T` stored in any table. If you would like to restrict the references to point only to objects stored in the `DEPARTMENTS` table, you could do so by adding a scope constraint on the `DEPT` column as follows:

```
ALTER TABLE emp
  ADD (SCOPE FOR (dept) IS departments);
```

The above `ALTER TABLE` statement will succeed *only if* the `EMP` table is empty.

If you want the `REF` values in the `DEPT` column of `EMP` to also store the rowids, issue the following statement:

```
ALTER TABLE emp
  ADD (REF(dept) WITH ROWID);
```

**Add Partition Example** The following statement adds partition `JAN99` to tablespace `TSX`:

```
ALTER TABLE sales
  ADD PARTITION jan99 VALUES LESS THAN( '970201' )
  TABLESPACE tsx;
```

**Drop Partition Example** The following statement drops partition `DEC98`:

```
ALTER TABLE sales DROP PARTITION dec98;
```

**Exchange Partition Example** The following statement converts partition FEB97 to table SALES\_FEB97 without exchanging local index partitions with corresponding indexes on SALES\_FEB97 and without verifying that data in SALES\_FEB97 falls within the bounds of partition FEB97:

```
ALTER TABLE sales
  EXCHANGE PARTITION feb97 WITH TABLE sales_feb97
  WITHOUT VALIDATION;
```

**Modify Partition Examples** The following statement marks all the local index partitions corresponding to the NOV96 partition of the SALES table UNUSABLE:

```
ALTER TABLE sales MODIFY PARTITION nov96
  UNUSABLE LOCAL INDEXES;
```

The following statement rebuilds all the local index partitions that were marked UNUSABLE:

```
ALTER TABLE sales MODIFY PARTITION jan97
  REBUILD UNUSABLE LOCAL INDEXES;
```

The following statement changes MAXEXTENTS and logging attribute for partition BRANCH\_NY:

```
ALTER TABLE branch MODIFY PARTITION branch_ny
  STORAGE (MAXEXTENTS 75) LOGGING;
```

**Move Partition Example** The following statement moves partition DEPOT2 to tablespace TS094:

```
ALTER TABLE parts
  MOVE PARTITION depot2 TABLESPACE ts094 NOLOGGING;
```

**Rename Partition Examples** The following statement renames a table:

```
ALTER TABLE emp RENAME TO employee;
```

In the following statement, partition EMP3 is renamed:

```
ALTER TABLE employee RENAME PARTITION emp3 TO employee3;
```

**Split Partition Example** The following statement splits the old partition DEPOT4, creating two new partitions, naming one DEPOT9 and reusing the name of the old partition for the other:

```
ALTER TABLE parts
  SPLIT PARTITION depot4 AT ( '40-001' )
  INTO ( PARTITION depot4 TABLESPACE ts009 STORAGE (MINEXTENTS 2),
        PARTITION depot9 TABLESPACE ts010 )
  PARALLEL (10);
```

**Truncate Partition Example** The following statement deletes all the data in the SYS\_P017 partition and deallocates the freed space:

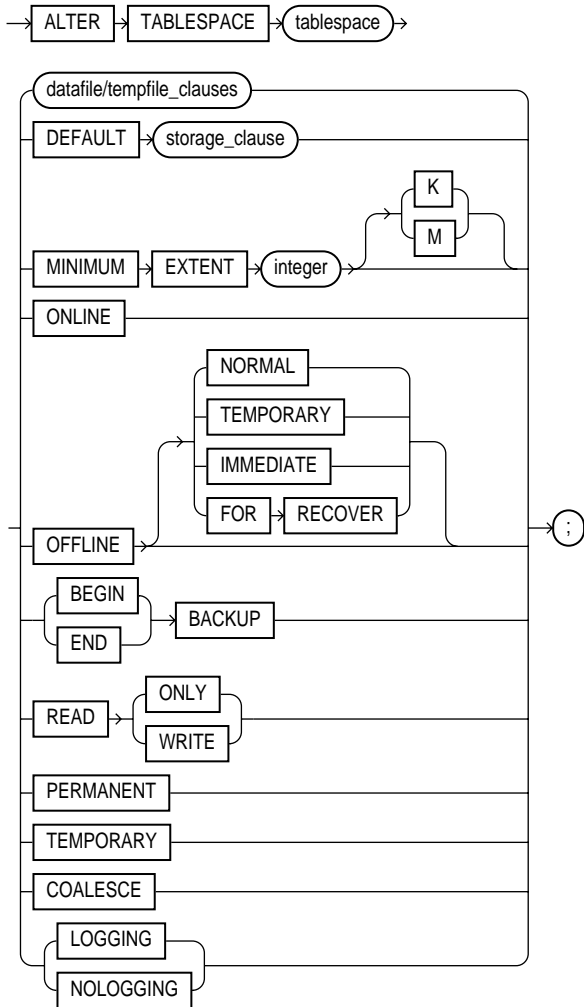
```
ALTER TABLE deliveries
  TRUNCATE PARTITION sys_p017 DROP STORAGE;
```


**Additional Examples** For examples of defining integrity constraints with the ALTER TABLE statement, see the "[constraint\\_clause](#)" on page 7-233.

For examples of changing the value of a table's storage parameters, see the "[storage\\_clause](#)" on page 7-605.



# ALTER TABLESPACE

## Syntax



**datafile/tempfile\_clauses::=**

**filespec:** See "filespec" on page 7-516.

**autoextend\_clause::=****maxsize\_clause::=**

**storage\_clause:** See "storage\_clause" on page 7-605.

**Purpose**

To alter an existing tablespace or one or more of its datafiles or tempfiles.

For information on creating a tablespace, see "[CREATE TABLESPACE](#)" on page 4-419.

## Prerequisites

If you have `ALTER TABLESPACE` system privilege, you can perform any of this statement's operations. If you have `MANAGE TABLESPACE` system privilege, you can only perform the following operations:

- take the tablespace online or offline
- begin or end a backup
- make the tablespace read-only or read-write

Before you can make a tablespace read-only, the following conditions must be met:

- The tablespace must be online.
- The tablespace must not contain any active rollback segments. For this reason, the `SYSTEM` tablespace can never be made read-only, because it contains the `SYSTEM` rollback segment. Additionally, because the rollback segments of a read-only tablespace are not accessible, Oracle recommends that you drop the rollback segments before you make a tablespace read-only.
- The tablespace must not be involved in an open backup, because the end of a backup updates the header file of all datafiles in the tablespace.

Performing this function in restricted mode may help you meet these restrictions, because only users with `RESTRICTED SESSION` system privilege can be logged on.

## Keywords and Parameters

---

<i>tablespace</i>	is the name of the tablespace to be altered.
	<b>Note:</b> For locally managed temporary tablespaces, the only clause you can specify in this statement in the <code>ADD</code> clause.
LOGGING   NOLOGGING	specifies the default logging attribute of all tables, indexes, and partitions within the tablespace. The tablespace-level logging attribute can be overridden by logging specifications at the table, index, and partition levels.  When an existing tablespace logging attribute is changed by an <code>ALTER TABLESPACE</code> statement, all tables, indexes, and partitions created <i>after</i> the statement will have the new default logging attribute (which you can still subsequently override). The logging attributes of existing objects are not changed.  Only the following operations support <code>NOLOGGING</code> mode: <ul style="list-style-type: none"><li>■ DML: direct-load <code>INSERT</code> (serial or parallel); Direct Loader (<code>SQL*Loader</code>)</li></ul>

- DDL: CREATE TABLE ... AS SELECT, CREATE INDEX, ALTER INDEX ... REBUILD, ALTER INDEX ... REBUILD PARTITION, ALTER INDEX ... SPLIT PARTITION, ALTER TABLE ... SPLIT PARTITION, ALTER TABLE ... MOVE PARTITION.

In NOLOGGING mode, data is modified with minimal logging (to mark new extents invalid and to record dictionary changes). When applied during media recovery, the extent invalidation records mark a range of blocks as logically corrupt, because the redo data is not logged. Therefore, if you cannot afford to lose the object, it is important to take a backup after the NOLOGGING operation.

*datafile/tempfile\_*  
*clauses* adds or modifies a datafile or tempfile.

ADD DATAFILE | TEMPFILE Adds to the tablespace a datafile or tempfile specified by *filespec* (see "filespec" on page 7-516).

You can add a datafile or tempfile to a locally managed tablespace that is online or to a dictionary managed tablespace that is online or offline. Be sure the file is not in use by another database.

**Note:** For locally managed temporary tablespaces, this is the only clause you can specify at any time.

RENAME DATAFILE renames one or more of the tablespace's datafiles. Take the tablespace offline before renaming the datafile. Each *filename* must fully specify a datafile using the conventions for filenames on your operating system.

This clause merely associates the tablespace with the new file rather than the old one. This clause does not actually change the name of the operating system file. You must change the name of the file through your operating system.

*autoextend\_*  
*clause* enables or disables the autoextending of the size of the datafile in the tablespace.

OFF disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in further ALTER TABLESPACE AUTOEXTEND statements.

ON enables autoextend.

NEXT specifies the size in bytes of the next increment of disk space to be allocated automatically to the datafile when more extents are required. Use K or M to specify this size in kilobytes or megabytes. The default is one data block.

*maxsize\_clause* specifies maximum disk space allowed for automatic extension of the datafile.

UNLIMITED sets no limit on allocating disk space to the datafile.

DEFAULT <i>storage_clause</i>	specifies the new default storage parameters for objects subsequently created in the tablespace. For a dictionary-managed temporary table, Oracle considers only the NEXT parameter of the <i>storage_clause</i> . See the " <a href="#">storage_clause</a> " on page 7-605. <b>Restriction:</b> You cannot specify this clause for a locally managed tablespace.
MINIMUM EXTENT <i>integer</i>	controls free space fragmentation in the tablespace by ensuring that every used or free extent size in a tablespace is at least as large as, and is a multiple of, <i>integer</i> . This clause is not relevant for a dictionary-managed temporary tablespace. <b>Restriction:</b> You cannot specify this clause for a locally managed tablespace. <b>See Also:</b> <i>Oracle8i Administrator's Guide</i> for more information about using MINIMUM EXTENT to control space fragmentation.
ONLINE	brings the tablespace online.
OFFLINE	takes the tablespace offline and prevents further access to its segments.
NORMAL	flushes all blocks in all datafiles in the tablespace out of the SGA. You need not perform media recovery on this tablespace before bringing it back online. This is the default.
TEMPORARY	performs a checkpoint for all online datafiles in the tablespace but does not ensure that all files can be written. Any offline files may require media recovery before you bring the tablespace back online.
IMMEDIATE	does not ensure that tablespace files are available and does not perform a checkpoint. You must perform media recovery on the tablespace before bringing it back online.
FOR RECOVER	takes the production database tablespaces in the recovery set offline for tablespace point-in-time recovery. For additional information see <i>Oracle8i Backup and Recovery Guide</i> .
<b>Suggestion:</b>	Before taking a tablespace offline for a long time, you may want to alter the tablespace allocation of any users who have been assigned the tablespace as either a default or temporary tablespace. When the tablespace is offline, these users cannot allocate space for objects or sort areas in the tablespace.
<b>See Also:</b>	"ALTER USER" on page 7-193.
BEGIN BACKUP	signifies that an open backup is to be performed on the datafiles that make up this tablespace. This clause does not prevent users from accessing the tablespace. You must use this clause before beginning an open backup. You cannot use this clause on a read-only tablespace. <hr/> <b>Note:</b> While the backup is in progress, you cannot take the tablespace offline normally, shut down the instance, or begin another backup of the tablespace. <hr/>
END BACKUP	signifies that an open backup of the tablespace is complete. Use this clause as soon as possible after completing an open backup. You cannot use this clause on a read-only tablespace.



---

If you forget to indicate the end of an online tablespace backup, and an instance failure or SHUTDOWN ABORT occurs, Oracle assumes that media recovery (possibly requiring archived redo log) is necessary at the next instance start up. To restart the database without media recovery, see *Oracle8i Administrator's Guide*.

READ ONLY	places the tablespace in <b>transition read-only mode</b> . In this state, existing transactions can complete (commit or roll back), but no further write operations (DML) are allowed to the tablespace except for rollback of existing transactions that previously modified blocks in the tablespace.  Once a tablespace is read only, you can copy its files to read-only media. You must then rename the datafiles in the control file to point to the new location by using the SQL statement ALTER DATABASE ... RENAME. See " <a href="#">ALTER DATABASE</a> " on page 7-6.  <b>See Also:</b> <i>Oracle8i Concepts</i> for more information on read-only tablespaces.
READ WRITE	signifies that write operations are allowed on a previously read-only tablespace.
PERMANENT	specifies that the tablespace is to be converted from a temporary to a permanent one. A permanent tablespace is one in which permanent database objects can be stored. This is the default when a tablespace is created.
TEMPORARY	specifies that the tablespace is to be converted from a permanent to a temporary one. A temporary tablespace is one in which no permanent database objects can be stored. Objects in a temporary tablespace persist only for the duration of the session.
COALESCE	for each datafile in the tablespace, coalesces all contiguous free extents into larger contiguous extents.

---

## Examples

**Backup Examples** The following statement signals to the database that a backup is about to begin:

```
ALTER TABLESPACE accounting
  BEGIN BACKUP;
```

The following statement signals to the database that the backup is finished:

```
ALTER TABLESPACE accounting
  END BACKUP;
```

**Moving and Renaming Example** This example moves and renames a datafile associated with the ACCOUNTING tablespace from 'DISKA:PAY1.DAT' to 'DISKB:RECEIVE1.DAT':

1. Take the tablespace offline using an ALTER TABLESPACE statement with the OFFLINE clause:

```
ALTER TABLESPACE accounting OFFLINE NORMAL;
```

2. Copy the file from 'DISKA:PAY1.DAT' to 'DISKB:RECEIVE1.DAT' using your operating system's commands.
3. Rename the datafile using the ALTER TABLESPACE statement with the RENAME DATAFILE clause:

```
ALTER TABLESPACE accounting
  RENAME DATAFILE 'diska:pay1.dbf'
  TO      'diskb:receive1.dbf';
```

4. Bring the tablespace back online using an ALTER TABLESPACE statement with the ONLINE clause:

```
ALTER TABLESPACE accounting ONLINE;
```

**Adding a Datafile Example** The following statement adds a datafile to the tablespace and changes the default logging attribute to NOLOGGING. When more space is needed, new extents of size 10 kilobytes will be added up to a maximum of 100 kilobytes:

```
ALTER TABLESPACE accounting NOLOGGING
  ADD DATAFILE 'disk3:pay3.dbf'
  SIZE 50K
  AUTOEXTEND ON
  NEXT 10K
  MAXSIZE 100K;
```


Altering a tablespace logging attribute has no effect on the logging attributes of the existing schema objects within the tablespace. The tablespace-level logging attribute can be overridden by logging specifications at the table, index, and partition levels.

**Changing Extent Allocation Example** The following statement changes the allocation of every extent of TABSPACE\_ST to a multiple of 128K:

```
ALTER TABLESPACE tabspace_st MINIMUM EXTENT 128K;
```

## ALTER TRIGGER

### Syntax



### Purpose

To enable, disable, or compile a database trigger. For information on creating a trigger, see ["CREATE TRIGGER"](#) on page 7-426. For information on dropping a trigger, see ["DROP TRIGGER"](#) on page 7-505.

---



---

**Note:** This statement does not change the declaration or definition of an existing trigger. To redeclare or redefine a trigger, use the `CREATE TRIGGER` statement with `OR REPLACE`.

---



---

### Prerequisites

The trigger must be in your own schema or you must have `ALTER ANY TRIGGER` system privilege.

In addition, to alter a trigger on `DATABASE`, you must have the `ADMINISTER DATABASE TRIGGER` system privilege.

**See Also:** ["CREATE TRIGGER"](#) on page 7-426 for more information on triggers based on `DATABASE`.

### Keywords and Parameters

<i>schema</i>	is the schema containing the trigger. If you omit <i>schema</i> , Oracle assumes the trigger is in your own schema.
<i>trigger</i>	is the name of the trigger to be altered.

ENABLE	enables the trigger. You can also use the <code>ENABLE ALL TRIGGERS</code> clause of <code>ALTER TABLE</code> to enable all triggers associated with a table. See " <a href="#">ALTER TABLE</a> " on page 7-123.
DISABLE	disables the trigger. You can also use the <code>DISABLE ALL TRIGGERS</code> clause of <code>ALTER TABLE</code> to disable all triggers associated with a table. See " <a href="#">ALTER TABLE</a> " on page 7-123.
COMPILE	explicitly compiles the trigger, whether it is valid or invalid. Explicit recompilation eliminates the need for implicit run-time recompilation and prevents associated run-time compilation errors and performance overhead.  Oracle first recompiles objects upon which the trigger depends, if any of these objects are invalid. If Oracle recompiles the trigger successfully, the trigger becomes valid.  If recompiling the trigger results in compilation errors, then Oracle returns an error and the trigger remains invalid. You can see the associated compiler error messages with the SQL*Plus command <code>SHOW ERRORS</code> . For information on debugging procedures, see <i>Oracle8i Application Developer's Guide - Fundamentals</i> . For information on how Oracle maintains dependencies among schema objects, including remote objects, see <i>Oracle8i Concepts</i> .
DEBUG	instructs the PL/SQL compiler to generate and store the code for use by the PL/SQL debugger. This clause can be used for normal triggers and for instead-of triggers.

---

## Examples

Consider a trigger named `REORDER` created on the `INVENTORY` table. The trigger is fired whenever an `UPDATE` statement reduces the number of a particular part on hand below the part's reorder point. The trigger inserts into a table of pending orders a row that contains the part number, a reorder quantity, and the current date.

When this trigger is created, Oracle enables it automatically. You can subsequently disable the trigger with the following statement:

```
ALTER TRIGGER reorder DISABLE;
```

When the trigger is disabled, Oracle does not fire the trigger when an `UPDATE` statement causes the part's inventory to fall below its reorder point.

After disabling the trigger, you can subsequently enable it with the following statement:

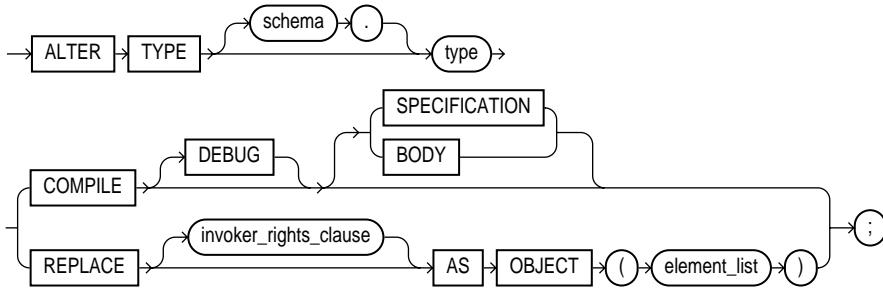
```
ALTER TRIGGER reorder ENABLE;
```

After you reenable the trigger, Oracle fires the trigger whenever a part's inventory falls below its reorder point as a result of an `UPDATE` statement. It is possible that a part's inventory falls below its reorder point while the trigger was disabled. In that


case, when you reenable the trigger, Oracle does not automatically fire the trigger for this part until another transaction further reduces the inventory.

# ALTER TYPE


## Syntax



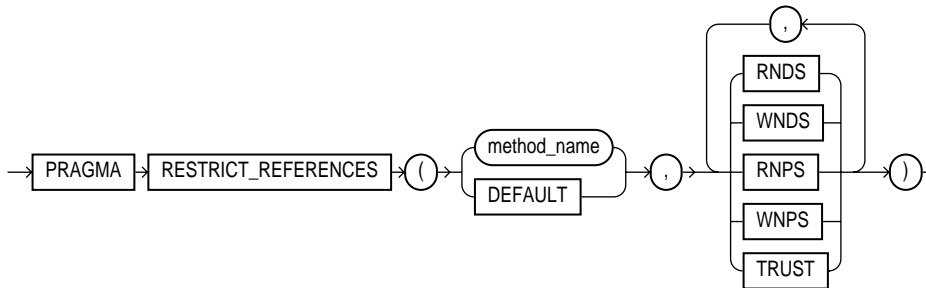
### element\_list::=



### invoker\_rights\_clause::=



**pragma\_clause::=**



## Purpose

To recompile the specification and/or body, or to change the specification of an object type by adding new object member subprogram specifications.

You cannot change the existing properties (attributes, member subprograms, map or order functions) of an object type, but you can add new member subprogram specifications.

## Prerequisites

The object type must be in your own schema and you must have `CREATE TYPE` or `CREATE ANY TYPE` system privilege, or you must have `ALTER ANY TYPE` system privileges.

## Keywords and Parameters

---

<i>schema</i>	is the schema that contains the type. If you omit <i>schema</i> , Oracle assumes the type is in your current schema.
<i>type</i>	is the name of an object type, a nested table type, or a rowid type.
COMPILE	compiles the object type specification and body. This is the default if neither <code>SPECIFICATION</code> nor <code>BODY</code> is specified.
	If recompiling the type results in compilation errors, then Oracle returns an error and the type remains invalid. You can see the associated compiler error messages with the SQL*Plus command <code>SHOW ERRORS</code> .
<code>SPECIFICATION</code>	compiles only the object type specification.
<code>BODY</code>	compiles only the object type body.

---

DEBUG	instructs the PL/SQL compiler to generate and store the code for use by the PL/SQL debugger.
REPLACE AS OBJECT	adds new member subprogram specifications. This clause is valid only for object types, not for nested table or varray types.
<i>invoker_rights_ clause</i>	<p>specifies whether the member functions and procedures of the object type execute with the privileges and in the schema of the user who owns the object type or with the privileges and in the schema of <code>CURRENT_USER</code>. This specification applies to the corresponding type body as well. (For information on how <code>CURRENT_USER</code> is determined, see <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i>.)</p> <p>This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the member functions and procedures of the type.</p> <p><b>Restriction:</b> You can specify this clause only for an object type, not for a nested table or varray type.</p> <p><b>See Also:</b> <i>PL/SQL User's Guide and Reference</i>.</p>
AUTHID CURRENT_USER	<p>specifies that the member functions and procedures of the object type execute with the privileges of <code>CURRENT_USER</code>. This clause creates an "invoker-rights type."</p> <p>This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of <code>CURRENT_USER</code>. External names in all other statements resolve in the schema in which the type resides.</p> <hr/> <p><b>Note:</b> You must specify this clause to maintain invoker-rights status for the type if you created it with this status. Otherwise the status will revert to definer rights.</p> <hr/>
AUTHID DEFINER	specifies that the member functions and procedures of the object type execute with the privileges of the owner of the schema in which the functions and procedures reside, and that external names resolve in the schema where the member functions and procedures reside. This is the default.
<i>attribute</i>	is an object attribute name. Attributes are data items with a name and a type specifier that form the structure of the object.
MEMBER   STATIC	<p>specifies a function or procedure subprogram associated with the object type which is referenced as an attribute. For a description of the difference between member and static methods, and for examples, see "<a href="#">CREATE TYPE</a>" on page 7-437. For information about overloading subprogram names within a package, see the <i>PL/SQL User's Guide and Reference</i>.</p> <p>You must specify a corresponding method body in the object type body for each procedure or function specification. See "<a href="#">CREATE TYPE BODY</a>" on page 7-447.</p>
<i>procedure_spec</i>	is the specification of a procedure subprogram.



---

	<i>function_spec</i>	is the specification of a function subprogram.
<i>pragma_clause</i>		is a compiler directive that denies member functions read/write access to database tables, packaged variables, or both, and thereby helps to avoid side effects. <b>See Also:</b> <i>Oracle8i Application Developer's Guide - Fundamentals</i> .
	<i>method</i>	is the name of the MEMBER function or procedure to which the pragma is being applied.
	DEFAULT	specifies that the pragma should be applied to all methods in the type for which a pragma has not been explicitly specified.
	WNDS	specifies the constraint <i>writes no database state</i> (does not modify database tables).
	WNPS	specifies the constraint <i>writes no package state</i> (does not modify packaged variables).
	RNDS	specifies the constraint <i>reads no database state</i> (does not query database tables).
	RNPS	specifies the constraint <i>reads no package state</i> (does not reference package variables).
	TRUST	specifies that the restrictions listed in the pragma are not actually to be enforced, but are simply trusted to be true.
MAP   ORDER MEMBER <i>function_spec</i>		
	MAP	specifies a member function (MAP method) that returns the relative position of a given instance in the ordering of all instances of the object. A map method is called implicitly and induces an ordering of object instances by mapping them to values of a predefined <i>scalar</i> type. Oracle uses the ordering for comparison operators and ORDER BY clauses.  If the argument to the map method is null, the map method returns null and the method is not invoked.  An object specification can contain only one map method, which must be a function. The result type must be a predefined SQL scalar type, and the map function can have no arguments other than the implicit SELF argument.
		<b>Note:</b> If <i>type_name</i> will be referenced in queries involving sorts (through ORDER BY, GROUP BY, DISTINCT, or UNION clauses) or joins, and you want those queries to be parallelized, you must specify a MAP member function.

**ORDER** specifies a member function (**ORDER** method) that takes an instance of an object as an explicit argument and the implicit **SELF** argument and returns either a negative, zero, or positive integer. The negative, zero, or positive indicates that the implicit **SELF** argument is less than, equal to, or greater than the explicit argument.

If either argument to the order method is null, the order method returns null and the method is not invoked.

When instances of the same object type definition are compared in an **ORDER BY** clause, the order method function is invoked.

An object specification can contain only one **ORDER** method, which must be a function having the return type **NUMBER**.

You can declare either a **MAP** method or an **ORDER** method, but not both. If you declare either method, you can compare object instances in SQL.

If you do not declare either method, you can compare object instances only for equality or inequality. Instances of the same type definition are equal only if each pair of their corresponding attributes is equal. No comparison method needs to be specified to determine the equality of two object types.

**See Also:** ["Object Values"](#) on page 2-35 for more information about object value comparisons.

---

## Examples

**Adding a Member Function** In the following example, member function **QTR** is added to the type definition of **DATA\_T**.

```
CREATE TYPE data_t AS OBJECT
  ( year NUMBER,
    MEMBER FUNCTION prod(invent NUMBER) RETURN NUMBER
  );

CREATE TYPE BODY data_t IS
  MEMBER FUNCTION prod (invent NUMBER) RETURN NUMBER IS
    BEGIN
      RETURN (year + invent);
    END;
END;

ALTER TYPE data_t REPLACE AS OBJECT
  ( year NUMBER,
    MEMBER FUNCTION prod(invent NUMBER) RETURN NUMBER,
    MEMBER FUNCTION qtr(der_qtr DATE) RETURN CHAR
  );
```

```

CREATE OR REPLACE TYPE BODY data_t IS
  MEMBER FUNCTION prod (invent NUMBER) RETURN NUMBER IS
  MEMBER FUNCTION qtr(der_qtr DATE) RETURN CHAR IS
    BEGIN
      RETURN (year + invent);
    END;
    BEGIN
      RETURN 'FIRST';
    END;
END;

```

**Recompiling a Type** The following example creates and then recompiles type LOAN\_T:

```

CREATE TYPE loan_t AS OBJECT
( loan_num      NUMBER,
  interest_rate  FLOAT,
  amount         FLOAT,
  start_date     DATE,
  end_date       DATE );

ALTER TYPE loan_t COMPILE;

```

**Recompiling a Type Body** The following example compiles the type body of LINK2.

```

CREATE TYPE link1 AS OBJECT
(a NUMBER);

CREATE TYPE link2 AS OBJECT
(a NUMBER,
 b link1,
 MEMBER FUNCTION p(c1 NUMBER) RETURN NUMBER);

CREATE TYPE BODY link2 AS
  MEMBER FUNCTION p(c1 NUMBER) RETURN NUMBER IS t13 link1;
  BEGIN t13 := link1(13);
        dbms_output.put_line(t13.a);
  RETURN 5;
  END;
END;

CREATE TYPE link3 AS OBJECT (a link2);
CREATE TYPE link4 AS OBJECT (a link3);

```

```
CREATE TYPE link5 AS OBJECT (a link4);
ALTER TYPE link2 COMPILE BODY;
```

**Recompiling a Type Specification** The following example compiles the type specification of LINK2.

```
CREATE TYPE link1 AS OBJECT
  (a NUMBER);

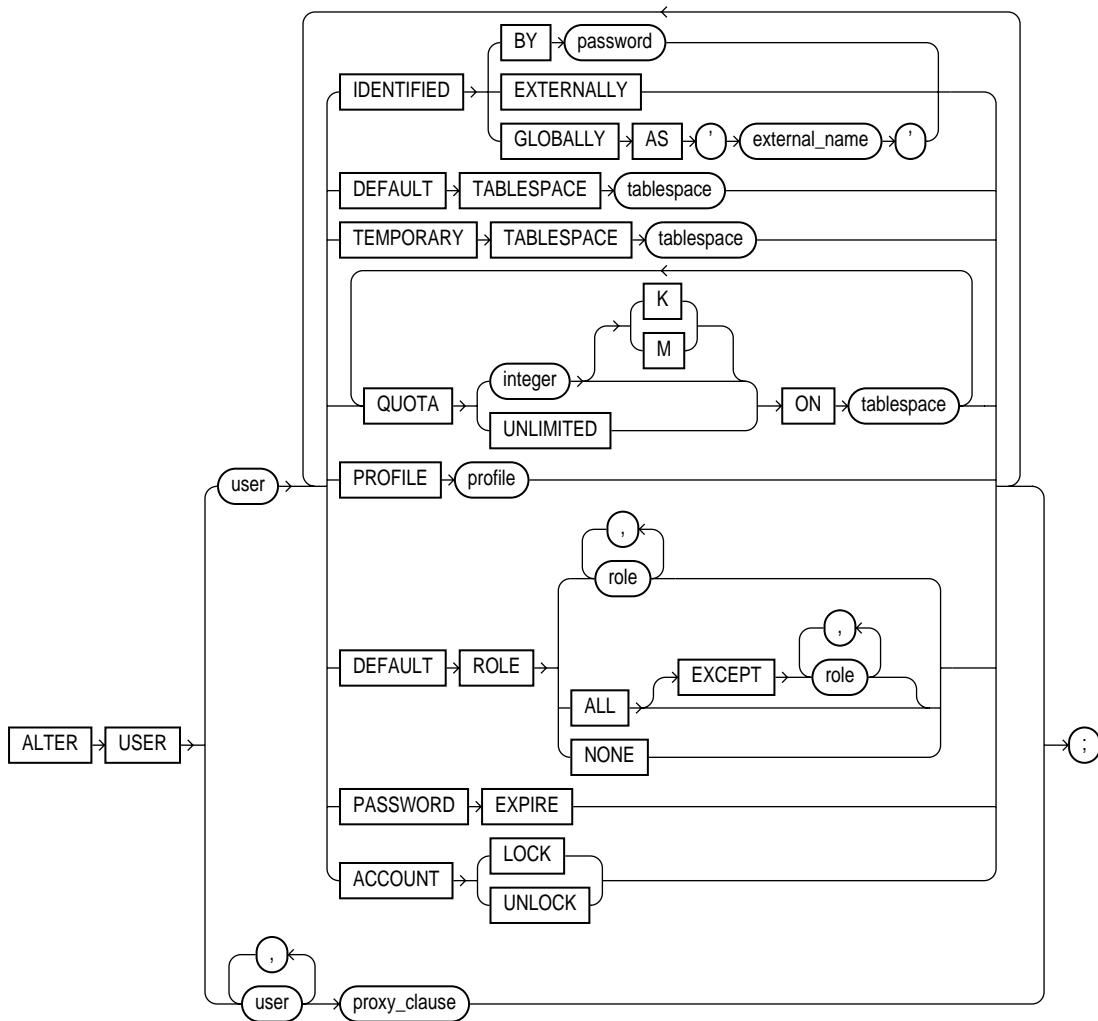
CREATE TYPE link2 AS OBJECT
  (a NUMBER,
   b link1,
   MEMBER FUNCTION p(c1 NUMBER) RETURN NUMBER);

CREATE TYPE BODY link2 AS
  MEMBER FUNCTION p(c1 NUMBER) RETURN NUMBER IS t14 link1;
  BEGIN t14 := link1(14);
    dbms_output.put_line(t14.a);
  RETURN 5;
  END;
END;

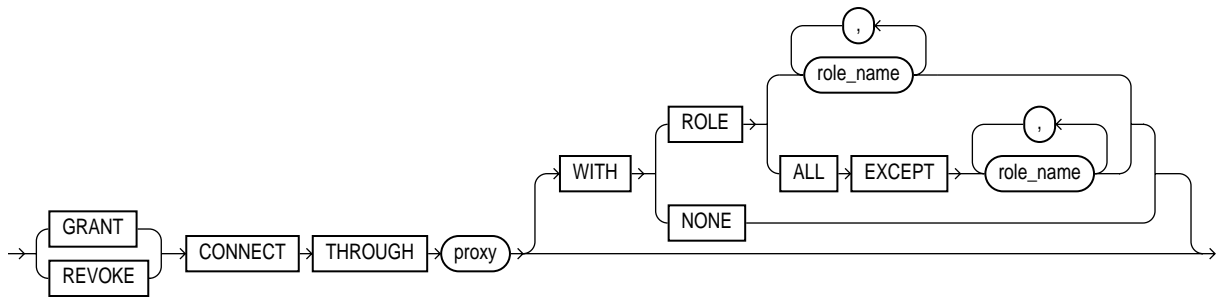
CREATE TYPE link3 AS OBJECT (a link2);
CREATE TYPE link4 AS OBJECT (a link3);
CREATE TYPE link5 AS OBJECT (a link4);
ALTER TYPE link2 COMPILE SPECIFICATION;
```

## ALTER USER

## Syntax



**proxy\_clause::=**



## Purpose

To change the authentication or database resource characteristics of a database user.

To permit a proxy server to connect as a client without authentication.

---



---

**Note:** ALTER USER syntax does not accept the old password. Therefore it neither authenticates using the old password nor checks the new password against the old before setting the new password. If these checks against the old password are important, use the OCIPasswordChange() call instead of ALTER USER. For more information, see *Oracle Call Interface Programmer's Guide*.

---



---

## Prerequisites

You must have the ALTER USER system privilege. However, you can change your own password without this privilege.

## Keywords and Parameters

The keywords and parameters shown below are unique to ALTER USER or have different functionality than they have in CREATE USER. All the remaining keywords and parameters in the ALTER USER statement have the same meaning as in the CREATE USER statement. For information on these keywords and parameters, see "CREATE USER" on page 7-451.

To assign limits on database resources to a user, see "CREATE PROFILE" on page 7-359.

---

<code>IDENTIFIED BY</code> <i>password</i>	<b>Note:</b> Oracle expects a different timestamp for each resetting of a particular password. If you reset one password multiple times within one second (for example, by cycling through a set of passwords using a script), Oracle may return an error message that the password cannot be reused. For this reason, Oracle Corporation recommends that you avoid using scripts to reset passwords.
<code>IDENTIFIED GLOBALLY AS</code>	indicates that a user must be authenticated by way of an LDAP V3 compliant directory service such as Oracle Internet Directory. (See also " <a href="#">CREATE USER</a> " on page 7-451.)  You can change a user's access verification method to <code>IDENTIFIED GLOBALLY AS 'external_name'</code> only if all external roles granted directly to the user are revoked.  You can change a user created as <code>IDENTIFIED GLOBALLY AS 'external_name'</code> to <code>IDENTIFIED BY password</code> or <code>IDENTIFIED EXTERNALLY</code> .
<code>DEFAULT ROLE</code>	can contain only roles that have been granted directly to the user with a <code>GRANT</code> statement. You cannot use the <code>DEFAULT ROLE</code> clause to enable: <ul style="list-style-type: none"> <li>■ roles not granted to the user</li> <li>■ roles granted through other roles</li> <li>■ roles managed by an external service (such as the operating system), or by the Oracle Internet Directory</li> </ul> <p>Oracle enables default roles at logon without requiring the user to specify their passwords.</p> <p><b>See Also:</b> "<a href="#">CREATE ROLE</a>" on page 4-365.</p>
<i>proxy_clause</i>	controls the ability of a <b>proxy</b> (an application or application server) to connect as the specified user and to activate all, some, or none of the user's roles.  <b>See Also:</b> <i>Oracle8i Concepts</i> for more information on proxies and their use of the database.
<code>GRANT</code>	allows the connection.
<code>REVOKE</code>	prohibits the connection.
<i>proxy</i>	identifies the proxy connecting to Oracle.
<code>WITH ROLE</code>	specifies the roles that the application is permitted to activate after it connects as the user. If you do not include this clause, Oracle activates all roles granted to the specified user automatically.
<i>role_name</i>	permits the proxy to connect as the specified user and to activate only the roles that are specified by <i>role_name</i> .
<code>ALL EXCEPT</code> <i>role_name</i>	permits the proxy to connect as the specified user and to activate all roles associated with that user except those specified by <i>role_name</i> .
<code>NONE</code>	permits the proxy to connect as the specified user, but prohibits the proxy from activating any of that user's roles after connecting.

---

## Examples

**General Examples** The following statement changes the user SCOTT's password to LION and default tablespace to the tablespace TSTEST:

```
ALTER USER scott
  IDENTIFIED BY lion
  DEFAULT TABLESPACE tstest;
```

The following statement assigns the CLERK profile to SCOTT:

```
ALTER USER scott
  PROFILE clerk;
```

In subsequent sessions, SCOTT is restricted by limits in the CLERK profile.

The following statement makes all roles granted directly to SCOTT default roles, except the AGENT role:

```
ALTER USER scott
  DEFAULT ROLE ALL EXCEPT agent;
```

At the beginning of SCOTT's next session, Oracle enables all roles granted directly to SCOTT except the AGENT role.

**Authentication Examples** The following statement changes user TOM's authentication mechanism:

```
ALTER USER tom IDENTIFIED GLOBALLY AS 'CN=tom,O=oracle,C=US';
```

The following statement causes user FRED's password to expire:

```
ALTER USER fred PASSWORD EXPIRE;
```

If you cause a database user's password to expire with `PASSWORD EXPIRE`, the user (or the DBA) must change the password before attempting to log in to the database following the expiration. However, tools such as SQL\*Plus allow you to change the password on the first attempted login following the expiration.

**Proxy Examples** The following statement permits the proxy user APPSERVER1 to connect as the user JANE. It also allows APPSERVER1 to activate the role INVENTORY:

```
ALTER USER jane GRANT CONNECT THROUGH appserver1 WITH ROLE inventory;
```




The following statement takes away the right of proxy user APPSERVER1 to connect as the user JANE:

```
ALTER USER jane REVOKE CONNECT THROUGH appserver1;
```

---

## ALTER VIEW

### Syntax



### Purpose

To explicitly recompile a view that is invalid. Explicit recompilation allows you to locate recompilation errors before run time. You may want to recompile a view explicitly after altering one of its base tables to ensure that the alteration does not affect the view or other objects that depend on it.

When you issue an `ALTER VIEW` statement, Oracle recompiles the view regardless of whether it is valid or invalid. Oracle also invalidates any local objects that depend on the view.

---

#### Notes:

- This statement does not change the definition of an existing view. To redefine a view, you must use `CREATE VIEW` with `OR REPLACE`. See "[CREATE VIEW](#)" on page 7-456.
- If you alter a view that is referenced by one or more materialized views, those materialized views are invalidated. Invalid materialized views cannot be used by query rewrite and cannot be refreshed. To revalidate an invalid materialized view, see "[ALTER MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-47. For information on materialized views in general, see *Oracle8i Data Warehousing Guide*.

**See Also:** *Oracle8i Concepts* for more about dependencies among schema objects.

---

### Prerequisites

The view must be in your own schema or you must have `ALTER ANY TABLE` system privilege.

---

## Keywords and Parameters

---

<i>schema</i>	is the schema containing the view. If you omit <i>schema</i> , Oracle assumes the view is in your own schema.
<i>view</i>	is the name of the view to be recompiled.
COMPILE	causes Oracle to recompile the view. The COMPILE keyword is required.

---

## Example

To recompile the view `CUSTOMER_VIEW`, issue the following statement:

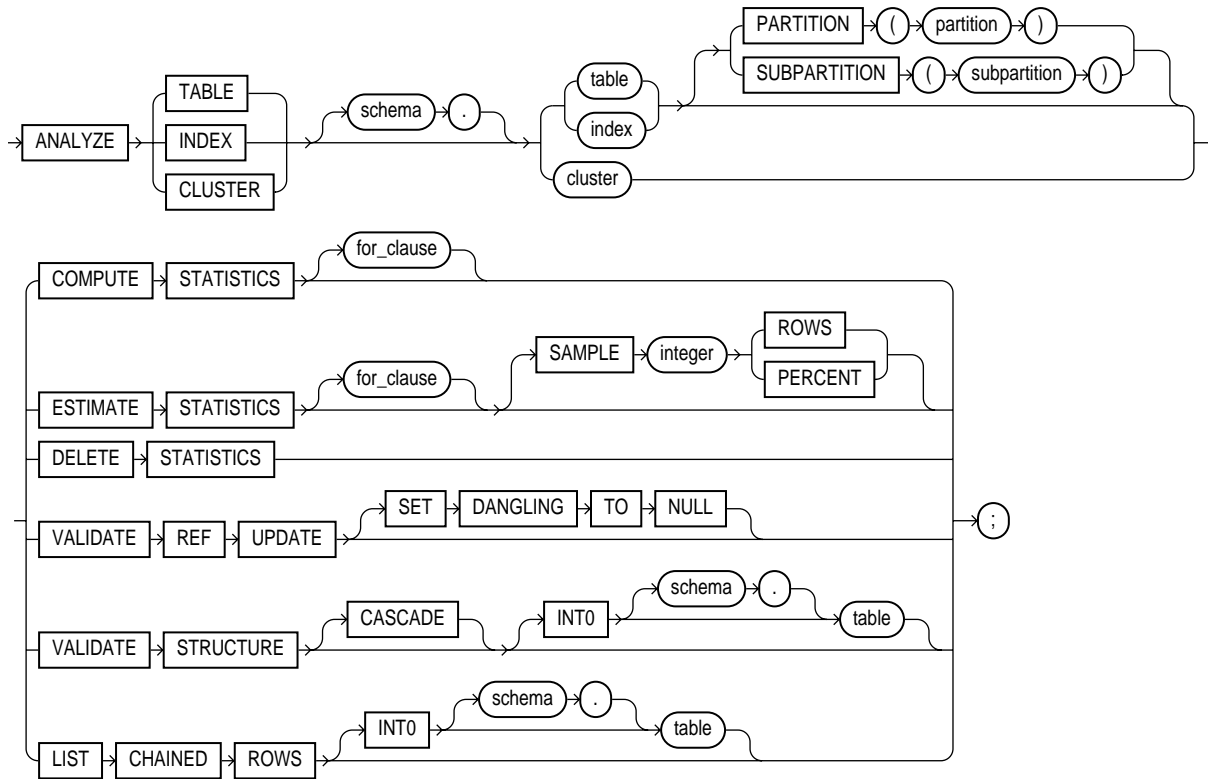
```
ALTER VIEW customer_view  
    COMPILE;
```

If Oracle encounters no compilation errors while recompiling `CUSTOMER_VIEW`, `CUSTOMER_VIEW` becomes valid. If recompiling results in compilation errors, Oracle returns an error and `CUSTOMER_VIEW` remains invalid.

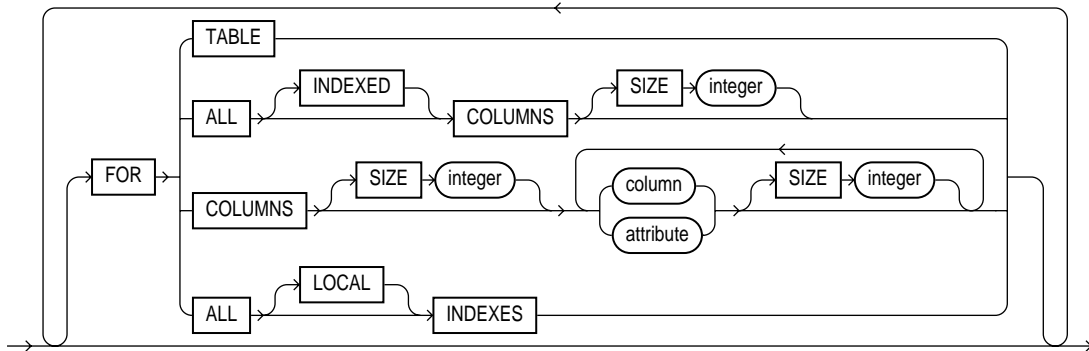
Oracle also invalidates all dependent objects. These objects include any procedures, functions, package bodies, and views that reference `CUSTOMER_VIEW`. If you subsequently reference one of these objects without first explicitly recompiling it, Oracle recompiles it implicitly at run time.

# ANALYZE

## Syntax



`for_clause::=`



## Purpose

This statement lets you

- Collect or delete statistics about an index or index partition, table or table partition, index-organized table, cluster, or scalar object attribute.
- Validate the structure of an index or index partition, table or table partition, index-organized table, cluster, or object reference (REF).
- Identify migrated and chained rows of a table or cluster.

For most statistics collection purposes, Oracle Corporation recommends that you use the `DBMS_STATS` package. That package lets you collect statistics in parallel, collect global statistics for partitioned objects, and fine tune your statistics collection in other ways. See *Oracle8i Supplied PL/SQL Packages Reference* for more information on this package.

However, you can use this statement for any of the purposes described in this section, and you **must** use this statement (rather than the `DBMS_STATS` package) for the following purposes:

- To use the `VALIDATE` or `LIST CHAINED ROWS` clauses
- To sample a number (rather than a percentage) of rows
- To collect statistics not used by the optimizer (such as information on freelist blocks)

## Prerequisites

The schema object to be analyzed must be local, and it must be in your own schema or you must have the `ANALYZE ANY` system privilege.

If you want to list chained rows of a table or cluster into a list table, the list table must be in your own schema, or you must have `INSERT` privilege on the list table, or you must have `INSERT ANY TABLE` system privilege.

If you want to validate a partitioned table, you must have `INSERT` privilege on the table into which you list analyzed rowids, or you must have `INSERT ANY TABLE` system privilege.

## Keywords and Parameters

---

*schema* is the schema containing the index, table, or cluster. If you omit *schema*, Oracle assumes the index, table, or cluster is in your own schema.

`INDEX index` identifies an index to be analyzed (if no *for\_clause* is used).

Oracle collects the following statistics for an index (statistics marked with an asterisk are always computed exactly):

- Depth of the index from its root block to its leaf blocks\*
- Number of leaf blocks
- Number of distinct index values
- Average number of leaf blocks per index value
- Average number of data blocks per index value (for an index on a table)
- Clustering factor (how well ordered the rows are about the indexed values)

Index statistics appear in the data dictionary views `USER_INDEXES`, `ALL_INDEXES`, and `DBA_INDEXES`.

For a **domain index**, this statement invokes the user-defined statistics collection function specified in the statistics type associated with the index (see "[ASSOCIATE STATISTICS](#)" on page 7-210). If no statistics type is associated with the domain index, the statistics type associated with its indextype is used. If no statistics type exists for either the index or its indextype, no user-defined statistics are collected. User-defined index statistics appear in the data dictionary views `USER_USTATS`, `ALL_USTATS`, and `DBA_USTATS`.

**Restriction:** You cannot analyze a domain index that is marked `LOADING` or `FAILED`.

**See Also:** "[CREATE INDEX](#)" on page 7-291 for more information on domain indexes.

---

TABLE <i>table</i>	<p>identifies a table to be analyzed. When you collect statistics for a table, Oracle also automatically collects the statistics for each of the table's indexes and domain indexes, provided that no <i>for_clauses</i> are used.</p> <p>When you analyze a table, Oracle collects statistics about expressions occurring in any function-based indexes as well. Therefore, be sure to create function-based indexes on the table before analyzing the table.</p> <p><b>See Also:</b> "<a href="#">CREATE INDEX</a>" on page 7-291 for more information about function-based indexes.</p> <p>When analyzing a table, Oracle skips all domain indexes marked <code>LOADING</code> or <code>FAILED</code>. Table statistics, including the status of domain indexes, appear in the data dictionary views <code>USER_TABLES</code>, <code>ALL_TABLES</code>, and <code>DBA_TABLES</code>.</p> <p>Oracle collects the following statistics for a table (statistics marked with an asterisk are always computed exactly):</p> <ul style="list-style-type: none"> <li>■ Number of rows</li> <li>■ * Number of data blocks below the high water mark (that is, the number of data blocks that have been formatted to receive data, regardless whether they currently contain data or are empty)</li> <li>■ * Number of data blocks allocated to the table that have never been used</li> <li>■ Average available free space in each data block in bytes</li> <li>■ Number of chained rows</li> <li>■ Average row length, including the row's overhead, in bytes</li> </ul> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot use <code>ANALYZE</code> to collect statistics on data dictionary tables.</li> <li>■ You cannot use <code>ANALYZE</code> to collect default statistics on a temporary table. However, if you have created an association between one or more columns of a temporary table and a user-defined statistics type, you can use <code>ANALYZE</code> to collect the user-defined statistics on the temporary table. (The association must already exist.)</li> <li>■ You cannot compute or estimate statistics for the following column types: <code>REFs</code>, <code>varrays</code>, <code>nested tables</code>, <code>LOBs</code> (<code>LOBs</code> are not analyzed, they are skipped), <code>LONGs</code>, or <code>object types</code>. However, if a statistics type is associated with such a column, user-defined statistics are collected.</li> </ul> <p><b>See Also:</b> "<a href="#">ASSOCIATE STATISTICS</a>" on page 7-210.</p>
PARTITION   SUBPARTITION	<p>specifies that statistics will be gathered for <i>partition</i> or <i>subpartition</i>. You cannot use this clause when analyzing clusters.</p> <p>If you specify <code>PARTITION</code> and <i>table</i> is composite-partitioned, Oracle analyzes all the subpartitions within the specified partition.</p>

---

**CLUSTER** *cluster* identifies a cluster to be analyzed. When you collect statistics for a cluster, Oracle also automatically collects the statistics for all the cluster's tables and all their indexes, including the cluster index.

- For an indexed cluster, Oracle collects the average number of data blocks taken up by a single cluster key value and all of its rows.
- For a hash cluster, Oracle collects the average number of data blocks taken up by a single hash key value and all of its rows.

These statistics appear in the data dictionary views `USER_CLUSTERS` and `DBA_CLUSTERS`.

**COMPUTE STATISTICS** computes exact statistics about the analyzed object and stores them in the data dictionary. When you analyze a table, both table and column statistics are collected.

**ESTIMATE STATISTICS** estimates statistics about the analyzed object and stores them in the data dictionary.

Both computed and estimated statistics are used by the Oracle optimizer to choose the execution plan for SQL statements that access analyzed objects. These statistics may also be useful to application developers who write such statements. For information on how these statistics are used, see *Oracle8i Designing and Tuning for Performance*.

**SAMPLE** *integer* specifies the amount of data from the analyzed object Oracle samples to estimate statistics. If you omit this parameter, Oracle samples 1064 rows.

The default sample value is adequate for tables up to a few thousand rows. If your tables are larger, specify a higher value for `SAMPLE`. If you specify more than half of the data, Oracle reads all the data and computes the statistics.

**ROWS** causes Oracle to sample *integer* rows of the table or cluster or *integer* entries from the index. The integer must be at least 1.

**PERCENT** causes Oracle to sample *integer* percent of the rows from the table or cluster or *integer* percent of the index entries. The integer can range from 1 to 99.

*for\_clause* specifies whether an entire table or index, or just particular columns, will be analyzed. The following clauses apply only to the `ANALYZE TABLE` version of this statement:

**FOR TABLE** restricts the statistics collected to only table statistics rather than table and column statistics.

**FOR COLUMNS** restricts the statistics collected to only column statistics for the specified columns and scalar object attributes, rather than for all columns and attributes; *attribute* specifies the qualified column name of an item in an object.

**FOR ALL COLUMNS** collects column statistics for all columns and scalar object attributes.



---

`FOR ALL INDEXED COLUMNS` collects column statistics for all indexed columns in the table.

Column statistics can be based on the entire column or can use a histogram by specifying `SIZE` (see below).

Oracle collects the following column statistics:

- Number of distinct values in the column as a whole
- Maximum and minimum values in each band

See Also: *Oracle8i Designing and Tuning for Performance* and "[Histogram Examples](#)" on page 7-208 for more information on histograms.

Column statistics appear in the data dictionary views `USER_TAB_COLUMNS`, `ALL_TAB_COLUMNS`, and `DBA_TAB_COLUMNS`. Histograms appear in the data dictionary views `USER_TAB_HISTOGRAMS`, `DBA_TAB_HISTOGRAMS`, and `ALL_TAB_HISTOGRAMS`; `USER_PART_HISTOGRAMS`, `DBA_PART_HISTOGRAMS`, and `ALL_PART_HISTOGRAMS`; and `USER_SUBPART_HISTOGRAMS`, `DBA_SUBPART_HISTOGRAMS`, and `ALL_SUBPART_HISTOGRAMS`.

**Note:** The `MAXVALUE` and `MINVALUE` columns of `USER_`, `DBA_`, and `ALL_TAB_COLUMNS` have a length of 32 bytes. If you analyze columns with a length >32 bytes, and if the columns are padded with leading blanks, Oracle may take into account only the leading blanks and return unexpected statistics.

If a user-defined statistics type has been associated with any columns, the *for\_clause* collects user-defined statistics using that statistics type. If no statistics type is associated with a column, Oracle checks to see if any statistics type has been associated with the type of the column, and uses that statistics type. If no statistics type has been associated with either the column or its user-defined type, no user-defined statistics are collected. User-defined column statistics appear in the data dictionary views `USER_USTATS`, `ALL_USTATS`, and `DBA_USTATS`.

If you want to collect statistics on both the table as a whole and on one or more columns, be sure to generate the statistics for the table first, and then for the columns. Otherwise, the table-only `ANALYZE` will overwrite the histograms generated by the column `ANALYZE`. For example, issue the following statements:

```
ANALYZE TABLE emp ESTIMATE STATISTICS;
ANALYZE TABLE emp ESTIMATE STATISTICS FOR ALL COLUMNS;
```

`FOR ALL INDEXES` specifies that all indexes associated with the table will be analyzed.

`FOR ALL LOCAL INDEXES` specifies that all local index partitions are analyzed. You must specify the keyword `LOCAL` if the `PARTITION` clause and `INDEX` are specified.

`SIZE` specifies the maximum number of partitions in the histogram. The default value is 75, minimum value is 1, and maximum value is 254.

DELETE  
STATISTICS

deletes any statistics about the analyzed object that are currently stored in the data dictionary. Use this statement when you no longer want Oracle to use the statistics.

When you use this clause on a table, Oracle also automatically removes statistics for all the table's indexes. When you use this clause on a cluster, Oracle also automatically removes statistics for all the cluster's tables and all their indexes, including the cluster index.

If user-defined column or index statistics were collected for an object, Oracle also removes the user-defined statistics by invoking the statistics deletion function specified in the statistics type that was used to collect the statistics.

VALIDATE REF  
UPDATE

validates the REFs in the specified table, checks the rowid portion in each REF, compares it with the true rowid, and corrects, if necessary. You can use this clause only when analyzing a table.

SET DANGLING  
TO NULL sets to NULL any REFs (whether or not scoped) in the specified table that are found to point to an invalid or nonexistent object.

---

**Note:** If the owner of the table does not have `SELECT` object privilege on the referenced objects, Oracle will consider them invalid and set them to `NULL`. Subsequently these REFs will not be available in a query, even if it is issued by user with appropriate privileges on the objects.

---

VALIDATE  
STRUCTURE

validates the structure of the analyzed object. The statistics collected by this clause are not used by the Oracle optimizer, as are statistics collected by the `COMPUTE STATISTICS` and `ESTIMATE STATISTICS` clauses.

- For a table, Oracle verifies the integrity of each of the table's data blocks and rows.
- For a cluster, Oracle automatically validates the structure of the cluster's tables.
- For a partitioned table, Oracle also verifies that the row belongs to the correct partition. If the row does not collate correctly, the rowid is inserted into the `INVALID_ROWS` table.
- For a temporary table, Oracle validates the structure of the table and its indexes during the current session.
- For an index, Oracle verifies the integrity of each data block in the index and checks for block corruption. This clause does not confirm that each row in the table has an index entry or that each index entry points to a row in the table. You can perform these operations by validating the structure of the table with the `CASCADE` clause.

Oracle stores statistics about the index in the data dictionary views `INDEX_STATS` and `INDEX_HISTOGRAM`, which are described in *Oracle8i Reference*.

Validating the structure of an object prevents `SELECT`, `INSERT`, `UPDATE`, and `DELETE` statements from concurrently accessing the object. Therefore, do not use this clause on the tables, clusters, and indexes of your production applications during periods of high database activity.

If Oracle encounters corruption in the structure of the object, an error message is returned to you. In this case, drop and re-create the object.

---

INTO	specifies a table into which Oracle lists the rowids of the partitions whose rows do not collate correctly. If you omit <i>schema</i> , Oracle assumes the list is in your own schema. If you omit this clause altogether, Oracle assumes that the table is named <code>INVALID_ROWS</code> . The SQL script used to create this table is <code>UTLVALID.SQL</code> .
CASCADE	<p>validates the structure of the indexes associated with the table or cluster. If you use this clause when validating a table, Oracle also validates the table's indexes. If you use this clause when validating a cluster, Oracle also validates all the clustered tables' indexes, including the cluster index.</p> <p>If you use this clause to validate an enabled (but previously disabled) function-based index, validation errors may result. In this case, you must rebuild the index.</p>
LIST CHAINED ROWS	<p>identifies migrated and chained rows of the analyzed table or cluster. You cannot use this clause when analyzing an index.</p> <p>INTO</p> <p>specifies a table into which Oracle lists the migrated and chained rows. If you omit <i>schema</i>, Oracle assumes the list table is in your own schema. If you omit this clause altogether, Oracle assumes that the table is named <code>CHAINED_ROWS</code>. The list table must be on your local database.</p> <p>You can create the <code>CHAINED_ROWS</code> table using one of these scripts:</p> <ul style="list-style-type: none"> <li>■ <code>UTLCHAIN.SQL</code> uses physical rowids. Therefore it can accommodate rows from conventional tables but not from index-organized tables. (See the Note that follows.)</li> <li>■ <code>UTLCHN1.SQL</code> uses universal rowids, so it can accommodate rows from both conventional and index-organized tables.</li> </ul> <p>If you create your own chained-rows table, it must follow the format prescribed by one of these two scripts. See <i>Oracle8i Migration</i> for compatibility issues related to the use of these scripts.</p> <hr/> <p><b>Note:</b> If you are analyzing index-organized tables based on primary keys (rather than universal rowids), you must create a separate chained-rows table for each index-organized table to accommodate its primary-key storage. Use the SQL scripts <code>DBMSIOTC.SQL</code> and <code>PRVTIOTC.PLB</code> to define the <code>BUILD_CHAIN_ROWS_TABLE</code> procedure, and then execute this procedure to create an <code>IOT_CHAINED_ROWS</code> table for each such index-organized table.</p> <p>For information on the SQL scripts, see the <code>DBMS_IOT</code> package in <i>Oracle8i Supplied PL/SQL Packages Reference</i>. For information on eliminating migrated and chained rows, see <i>Oracle8i Designing and Tuning for Performance</i>.</p>

---

## Examples

**Analyzing a Cluster** The following statement estimates statistics for the `CUST_HISTORY` table and all of its indexes:

```
ANALYZE TABLE cust_history
ESTIMATE STATISTICS;
```

**Deleting Statistics** The following statement deletes statistics about the CUST\_HISTORY table and all its indexes from the data dictionary:

```
ANALYZE TABLE cust_history
DELETE STATISTICS;
```

**Histogram Examples** The following statement creates a 10-band histogram on the SAL column of the EMP table:

```
ANALYZE TABLE emp
COMPUTE STATISTICS FOR COLUMNS sal SIZE 10;
```

You can also collect histograms for a single partition of a table. The following statement analyzes the EMP table partition P1:

```
ANALYZE TABLE emp PARTITION (p1) COMPUTE STATISTICS;
```

**Index Example** The following statement validates the structure of the index PARTS\_INDEX:

```
ANALYZE INDEX parts_index VALIDATE STRUCTURE;
```

**Table Examples** The following statement analyzes the EMP table and all of its indexes:

```
ANALYZE TABLE emp VALIDATE STRUCTURE CASCADE;
```

For a table, the VALIDATE REF UPDATE clause verifies the REFS in the specified table, checks the rowid portion of each REF, and then compares it with the true rowid. If the result is an incorrect rowid, the REF is updated so that the rowid portion is correct.

The following statement validates the REFS in the EMP table:

```
ANALYZE TABLE emp
VALIDATE REF UPDATE;
```

**Cluster Example** The following statement analyzes the ORDER\_CUSTS cluster, all of its tables, and all of their indexes, including the cluster index:

```
ANALYZE CLUSTER order_custs
VALIDATE STRUCTURE CASCADE;
```

**Chained Rows Example** The following statement collects information about all the chained rows of the table ORDER\_HIST:

```
ANALYZE TABLE order_hist
    LIST CHAINED ROWS INTO cr;
```

The preceding statement places the information into the table CR. You can then examine the rows with this query:

```
SELECT *
    FROM cr;
```


OWNER_NAME	TABLE_NAME	CLUSTER_NAME	HEAD_ROWID	TIMESTAMP
-----	-----	-----	-----	-----
SCOTT	ORDER_HIST		AAAAZzAABAAABrXAAA	15-MAR-96

**COMPUTE Example** The following statement calculates statistics for a scalar object attribute:


```
ANALYZE TABLE emp COMPUTE STATISTICS FOR COLUMNS addr.street;
```

# ASSOCIATE STATISTICS

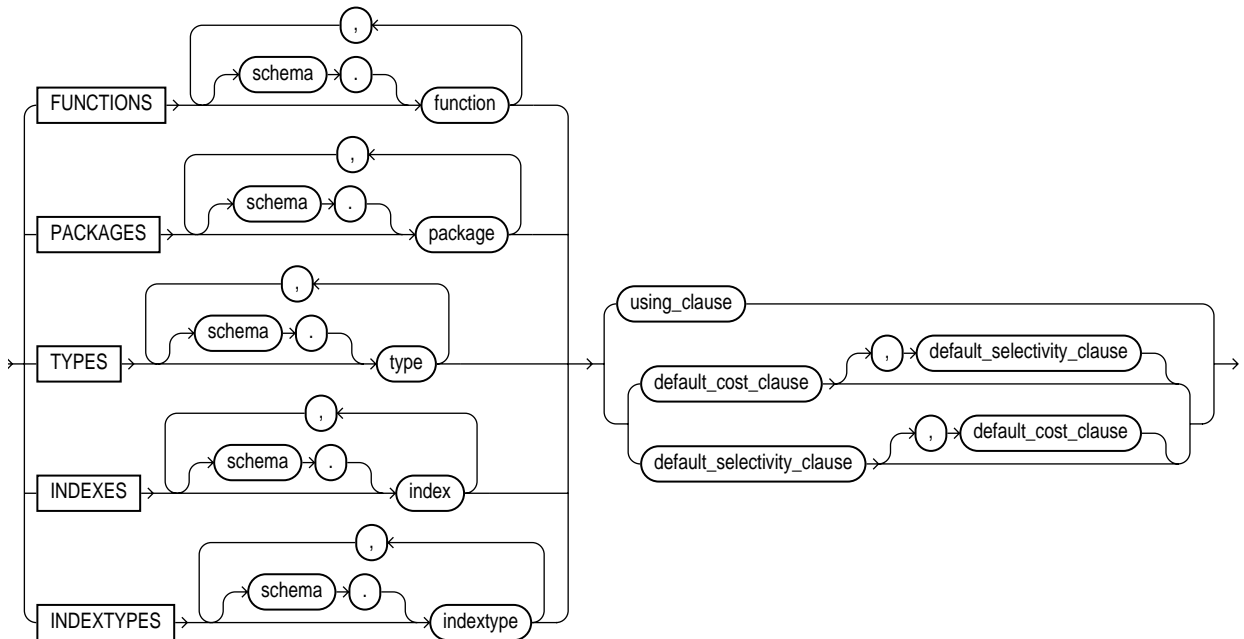
## Syntax





### column\_association::=



### function\_association::=



**using\_clause::=****default\_cost\_clause::=****default\_selectivity\_clause::=****Purpose**

To associate a statistics type (or default statistics) containing functions relevant to statistics collection, selectivity, or cost with one or more columns, standalone functions, packages, types, domain indexes, or indextypes.

For a listing of all current statistics type associations, refer to the `USER_ASSOCIATIONS` table. If you analyze the object with which you are associating statistics, you can also view the associations in the `USER_USTATS` table. For information on the order of precedence with which `ANALYZE` uses associations, see "[ANALYZE](#)" on page 7-200.

**Prerequisites**

To issue this statement, you must have the appropriate privileges to alter the base object (table, function, package, type, domain index, or indextype). In addition, unless you are associating only default statistics, you must have execute privilege on the statistics type. The statistics type must already have been defined. For information on defining types, see "[CREATE TYPE](#)" on page 7-437.

**Keywords and Parameters**


---

<i>column_</i> <i>association</i>	specifies a list of one or more table columns. If you do not specify <i>schema</i> , Oracle assumes the table is in your own schema.
--------------------------------------	--

<i>function_association</i>	<p>specifies a list of one or more standalone functions, packages, user-defined datatypes, domain indexes, or indextypes. If you do not specify <i>schema</i>, Oracle assumes the object is in your own schema.</p> <ul style="list-style-type: none"><li>▪ <code>FUNCTIONS</code> refers only to standalone functions, not to method types or to built-in functions.</li><li>▪ <code>TYPES</code> refers only to user-defined types, not to internal SQL datatypes.</li></ul> <p><b>Restriction:</b> You cannot specify an object for which you have already defined an association. You must first disassociate the statistics from this object. See "<a href="#">DISASSOCIATE STATISTICS</a>" on page 7-470.</p>
<i>using_clause</i>	<p>specifies the statistics type being associated with columns, functions, packages, types, domain indexes, or indextypes. The <i>statistics_type</i> must already have been created.</p>
<i>default_cost_clause</i>	<p>specifies default costs for standalone functions, packages, types, domain indexes, or indextypes. If you specify this clause, you must include one number each for CPU cost, I/O cost, and network cost, in that order. Each cost is for a single execution of the function or method or for a single domain index access. Accepted values are integers of zero or greater.</p>
<i>default_selectivity_clause</i>	<p>specifies as a percent the default selectivity for predicates with standalone functions, types, packages, or user-defined operators. The <i>default_selectivity</i> must be a whole number between 0 and 100. Values outside this range are ignored.</p> <p><b>Restriction:</b> You cannot specify <code>DEFAULT SELECTIVITY</code> for domain indexes or indextypes.</p>

---

## Examples

**Standalone Function Example** This statement creates an association for a standalone function `FN` and causes the optimizer to call the appropriate cost function (if present) in the statistics type `STAT_FN`.

```
ASSOCIATE STATISTICS WITH FUNCTIONS fn USING stat_fn;
```

**Default Cost Example** This statement specifies that using the domain index `T_A` to implement a given predicate always has a CPU cost of 100, I/O of 5, and network cost of 0.

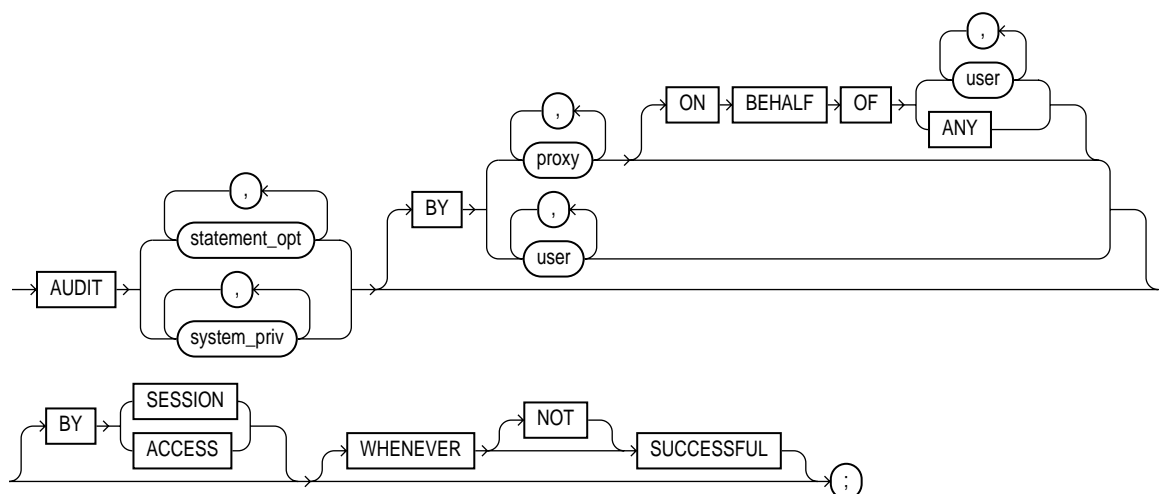
```
ASSOCIATE STATISTICS WITH INDEXES t_a DEFAULT COST (100,5,0);
```

The optimizer will simply use these default costs instead of calling a cost function.



## AUDIT *sql\_statements*

### Syntax



### Purpose

To track the occurrence of specific SQL statements in subsequent user sessions. Auditing options specified by the `AUDIT sql_statements` statement apply only to subsequent sessions, not to current sessions.

To choose particular schema objects for auditing, see "[AUDIT schema\\_objects](#)" on page 7-221. For information on disabling auditing of SQL statements, see "[NOAUDIT sql\\_statements](#)" on page 7-550.

### Prerequisites

You must have `AUDIT SYSTEM` system privilege.

You must enable auditing by setting the initialization parameter `AUDIT_TRAIL` to `DB`. You can specify auditing options regardless of whether auditing is enabled. However, Oracle does not generate audit records until you enable auditing.

## Keywords and Parameters

---

<i>statement_opt</i>	<p>chooses specific SQL statements for auditing. For a list of these statement options and the SQL statements they audit, see <a href="#">Table 7-1</a> and <a href="#">Table 7-2</a>.</p> <p>For each audited operation, Oracle produces an audit record containing this information:</p> <ul style="list-style-type: none"> <li>■ user performing the operation</li> <li>■ type of operation</li> <li>■ object involved in the operation</li> <li>■ date and time of the operation</li> </ul> <p>Oracle writes audit records to the audit trail, which is a database table containing audit records. You can review database activity by examining the audit trail through data dictionary views. For information on these views, see the <i>Oracle8i Reference</i>.</p>
<i>system_priv</i>	<p>chooses SQL statements that are authorized by the specified system privilege for auditing. For a list of all system privileges and the SQL statements that they authorize, see <a href="#">Table 7-5</a>.</p> <p>Rather than specifying many individual system privileges, you can specify the roles CONNECT, RESOURCE, and DBA. Doing so is equivalent to auditing all of the system privileges granted to those roles. See "<a href="#">GRANT system_privileges_and_roles</a>" on page 7-519 for more information on these roles.</p> <p>Oracle also provides two shortcuts for specifying groups of system privileges and statement options at once:</p> <ul style="list-style-type: none"> <li>■ ALL is equivalent to specifying all statements options shown in <a href="#">Table 7-1</a> but not the additional statement options shown in <a href="#">Table 7-2</a>.</li> <li>■ ALL PRIVILEGES is equivalent to specifying all system privileges.</li> </ul> <hr/> <p><b>Note:</b> Oracle Corporation recommends that you specify individual system privileges and statement options for auditing rather than roles or shortcuts. The specific system privileges and statement options encompassed by roles and shortcuts change from one release to the next and may not be supported in future versions of Oracle.</p>
BY <i>user</i>	chooses only SQL statements issued by specified users for auditing. If you omit this clause, Oracle audits all users' statements.
BY <i>proxy</i>	<p>chooses for auditing only SQL statements issued by the specified proxy.</p> <p><b>See Also:</b> <i>Oracle8i Concepts</i> for more information on proxies and their use of the database.</p> <p>ON BEHALF OF specifies the user or users on whose behalf the proxy executes the specified statement.</p> <ul style="list-style-type: none"> <li>■ <i>user</i> specifies auditing of statements executed on behalf of a particular user.</li> <li>■ ANY specifies auditing of statements executed on behalf of any user.</li> </ul>
BY SESSION	causes Oracle to write a single record for all SQL statements of the same type issued in the same session.

---

BY ACCESS	causes Oracle to write one record for each audited statement.  If you specify statement options or system privileges that audit data definition language (DDL) statements, Oracle automatically audits by access regardless of whether you specify the BY SESSION clause or BY ACCESS clause.  For statement options and system privileges that audit SQL statements other than DDL, you can specify either BY SESSION or BY ACCESS. BY SESSION is the default.
WHENEVER SUCCESSFUL	chooses auditing only for statements that succeed.  NOT chooses auditing only for statements that fail or result in errors.  If you omit the WHENEVER SUCCESSFUL clause, Oracle audits SQL statements regardless of success or failure.

---

**Table 7–1 Statement Auditing Options for Database Objects**

Statement Option	SQL Statements and Operations
CLUSTER	CREATE CLUSTER AUDIT CLUSTER DROP CLUSTER TRUNCATE CLUSTER
CONTEXT	CREATE CONTEXT DROP CONTEXT
DATABASE LINK	CREATE DATABASE LINK DROP DATABASE LINK
DIMENSION	CREATE DIMENSION ALTER DIMENSION DROP DIMENSION
DIRECTORY	CREATE DIRECTORY DROP DIRECTORY
INDEX	CREATE INDEX ALTER INDEX DROP INDEX
NOT EXISTS	All SQL statements that fail because a specified object does not exist.

**Table 7-1 (Cont.) Statement Auditing Options for Database Objects**

<b>Statement Option</b>	<b>SQL Statements and Operations</b>
PROCEDURE <sup>a</sup>	CREATE FUNCTION CREATE LIBRARY CREATE PACKAGE CREATE PACKAGE BODY CREATE PROCEDURE DROP FUNCTION DROP LIBRARY DROP PACKAGE DROP PROCEDURE
PROFILE	CREATE PROFILE ALTER PROFILE DROP PROFILE
PUBLIC DATABASE LINK	CREATE PUBLIC DATABASE LINK DROP PUBLIC DATABASE LINK
PUBLIC SYNONYM	CREATE PUBLIC SYNONYM DROP PUBLIC SYNONYM
ROLE	CREATE ROLE ALTER ROLE DROP ROLE SET ROLE
ROLLBACK STATEMENT	CREATE ROLLBACK SEGMENT ALTER ROLLBACK SEGMENT DROP ROLLBACK SEGMENT
SEQUENCE	CREATE SEQUENCE DROP SEQUENCE
SESSION	Logons
SYNONYM	CREATE SYNONYM DROP SYNONYM

**Table 7–1 (Cont.) Statement Auditing Options for Database Objects**

<b>Statement Option</b>	<b>SQL Statements and Operations</b>
SYSTEM AUDIT	AUDIT <i>sql_statements</i> NOAUDIT <i>sql_statements</i>
SYSTEM GRANT	GRANT <i>system_privileges_and_roles</i> REVOKE <i>system_privileges_and_roles</i>
TABLE	CREATE TABLE DROP TABLE TRUNCATE TABLE COMMENT ON TABLE DELETE [FROM] <i>table</i>
TABLESPACE	CREATE TABLESPACE ALTER TABLESPACE DROP TABLESPACE
TRIGGER	CREATE TRIGGER ALTER TRIGGER <b>with ENABLE and DISABLE clauses</b> DROP TRIGGER ALTER TABLE <b>with ENABLE ALL TRIGGERS clause</b> <b>and DISABLE ALL TRIGGERS clause</b>
TYPE	CREATE TYPE CREATE TYPE BODY ALTER TYPE DROP TYPE DROP TYPE BODY
USER	CREATE USER ALTER USER DROP USER
VIEW	CREATE VIEW DROP VIEW

**Table 7-1 (Cont.) Statement Auditing Options for Database Objects**

Statement Option	SQL Statements and Operations
	<sup>a</sup> Java schema objects (sources, classes, and resources) are considered the same as procedures for purposes of auditing SQL statements.

**Table 7-2 Additional Statement Auditing Options for SQL Statements**

Statement Option	SQL Statements and Operations
ALTER SEQUENCE	ALTER SEQUENCE
ALTER TABLE	ALTER TABLE
COMMENT TABLE	COMMENT ON TABLE <i>table</i> , <i>view</i> , <i>materialized view</i> COMMENT ON COLUMN <i>table.column</i> , <i>view.column</i> , <i>materialized view.column</i>
DELETE TABLE	DELETE FROM <i>table</i> , <i>view</i>
EXECUTE PROCEDURE	CALL  Execution of any procedure or function or access to any variable, library, or cursor inside a package.
GRANT DIRECTORY	GRANT privilege ON directory REVOKE privilege ON directory
GRANT PROCEDURE	GRANT privilege ON procedure, function, package REVOKE privilege ON procedure, function, package
GRANT SEQUENCE	GRANT privilege ON sequence REVOKE privilege ON sequence
GRANT TABLE	GRANT privilege ON table, view, materialized view. REVOKE privilege ON table, view, materialized view
GRANT TYPE	GRANT privilege ON TYPE REVOKE privilege ON TYPE
INSERT TABLE	INSERT INTO <i>table</i> , <i>view</i>
LOCK TABLE	LOCK TABLE <i>table</i> , <i>view</i>

**Table 7–2 (Cont.) Additional Statement Auditing Options for SQL Statements**

Statement Option	SQL Statements and Operations
SELECT SEQUENCE	Any statement containing <code>sequence.CURRVAL</code> or <code>sequence.NEXTVAL</code>
SELECT TABLE	SELECT FROM <code>table</code> , view, materialized view
UPDATE TABLE	UPDATE <code>table</code> , view

## Examples

**Role Examples** To choose auditing for every SQL statement that creates, alters, drops, or sets a role, regardless of whether the statement completes successfully, issue the following statement:

```
AUDIT ROLE;
```

To choose auditing for every statement that successfully creates, alters, drops, or sets a role, issue the following statement:

```
AUDIT ROLE
    WHENEVER SUCCESSFUL;
```

To choose auditing for every CREATE ROLE, ALTER ROLE, DROP ROLE, or SET ROLE statement that results in an Oracle error, issue the following statement:

```
AUDIT ROLE
    WHENEVER NOT SUCCESSFUL;
```

**Query/Update Examples** To choose auditing for any statement that queries or updates any table, issue the following statement:

```
AUDIT SELECT TABLE, UPDATE TABLE;
```

To choose auditing for statements issued by the users SCOTT and BLAKE that query or update a table or view, issue the following statement:

```
AUDIT SELECT TABLE, UPDATE TABLE
    BY scott, blake;
```

**Delete Example** To choose auditing for statements issued using the DELETE ANY TABLE system privilege, issue the following statement:

```
AUDIT DELETE ANY TABLE;
```

**Directory Examples** To choose auditing for statements issued using the CREATE ANY DIRECTORY system privilege, issue the following statement:

```
AUDIT CREATE ANY DIRECTORY;
```

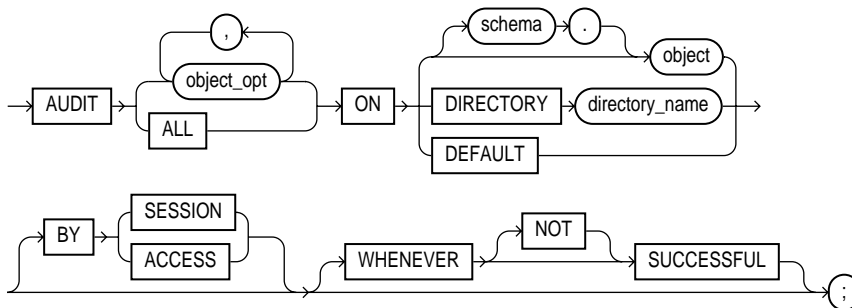
To choose auditing for CREATE DIRECTORY (and DROP DIRECTORY) statements that do not use the CREATE ANY DIRECTORY system privilege, issue the following statement:

```
AUDIT DIRECTORY;
```



## AUDIT *schema\_objects*

### Syntax



### Purpose

To track operations on a specific schema object. To choose particular SQL statements for auditing, see ["AUDIT sql\\_statements"](#) on page 7-213.

Auditing keeps track of operations performed by database users. Auditing options established by the `AUDIT schema_objects` statement apply to current sessions as well as to subsequent sessions. For information on discontinuing auditing operations, see ["NOAUDIT schema\\_objects"](#) on page 7-552.

### Prerequisites

The object you choose for auditing must be in your own schema or you must have `AUDIT ANY` system privilege. In addition, if the object you choose for auditing is a directory object, even if you created it, you must have `AUDIT ANY` system privilege.

### Keywords and Parameters

<i>object_opt</i>	specifies a particular operation for auditing. <a href="#">Table 7-3</a> shows each object option and the types of objects to which it applies. The name of each object option specifies a SQL statement to be audited. For example, if you choose to audit a table with the <code>ALTER</code> option, Oracle audits all <code>ALTER TABLE</code> statements issued against the table. If you choose to audit a sequence with the <code>SELECT</code> option, Oracle audits all statements that use any of the sequence's values.
-------------------	--

---

ALL	is a shortcut equivalent to specifying all object options applicable for the type of object. You can use this shortcut rather than explicitly specifying all options for an object.
<i>schema</i>	is the schema containing the object chosen for auditing. If you omit <i>schema</i> , Oracle assumes the object is in your own schema.
<i>object</i>	identifies the object chosen for auditing. The object must be a table, view, sequence, stored procedure, function, package, materialized view, or library.  You can also specify a synonym for a table, view, sequence, procedure, stored function, package, or materialized view.
ON DEFAULT	establishes the specified object options as default object options for subsequently created objects. Once you have established these default auditing options, any subsequently created object is automatically audited with those options. The default auditing options for a view are always the union of the auditing options for the view's base tables. You can see the current default auditing options by querying the ALL_DEF_AUDIT_OPTS data dictionary view.  If you change the default auditing options, the auditing options for previously created objects remain the same. You can change the auditing options for an existing object only by specifying the object in the ON clause of the AUDIT statement.
ON DIRECTORY <i>directory_name</i>	identifies the name of the directory chosen for auditing.
BY SESSION	causes Oracle to write a single record for all operations of the same type on the same object issued in the same session. This is the default.
BY ACCESS	causes Oracle to write one record for each audited operation.
WHENEVER SUCCESSFUL	chooses auditing only for SQL statements that complete successfully.  NOT chooses auditing only for statements that fail, or result in errors.  If you omit the WHENEVER SUCCESSFUL clause entirely, Oracle audits all SQL statements, regardless of success or failure.

---

**Table 7–3 Object Auditing Options**

Object Option	Table	View	Sequence	Procedure	Material- ized View / Snap- shot	Directory	Library	Object	
				Function Package <sup>a</sup>				Type	Context
ALTER	X		X		X			X	
AUDIT	X	X	X	X	X	X		X	X
COMMENT	X	X			X				
DELETE	X	X			X				
EXECUTE				X			X		
GRANT	X	X	X	X	X	X	X	X	X
INDEX	X				X				
INSERT	X	X			X				
LOCK	X	X			X				
READ						X			
RENAME	X	X		X	X				
SELECT	X	X	X		X				
UPDATE	X	X			X				

<sup>a</sup> Java schema objects (sources, classes, and resources) are considered the same as procedures, functions, and packages for purposes of auditing options.

## Examples

**Query Examples** To choose auditing for every SQL statement that queries the EMP table in the schema SCOTT, issue the following statement:

```
AUDIT SELECT
  ON scott.emp;
```

To choose auditing for every statement that successfully queries the EMP table in the schema SCOTT, issue the following statement:

```
AUDIT SELECT
  ON scott.emp
  WHENEVER SUCCESSFUL;
```

To choose auditing for every statement that queries the EMP table in the schema SCOTT and results in an Oracle error, issue the following statement:

```
AUDIT SELECT
  ON scott.emp
  WHENEVER NOT SUCCESSFUL;
```

**Insert/Update Example** To choose auditing for every statement that inserts or updates a row in the DEPT table in the schema BLAKE, issue the following statement:

```
AUDIT INSERT, UPDATE
  ON blake.dept;
```

**ALL Example** To choose auditing for every statement that performs any operation on the ORDER sequence in the schema ADAMS, issue the following statement:

```
AUDIT ALL
  ON adams.order;
```

The above statement uses the ALL shortcut to choose auditing for the following statements that operate on the sequence:

- ALTER SEQUENCE
- AUDIT
- GRANT
- any statement that accesses the sequence's values using the pseudocolumns CURRVAL or NEXTVAL

**READ Example** To choose auditing for every statement that reads files from the BFILE\_DIR1 directory, issue the following statement:

```
AUDIT READ ON DIRECTORY bfile_dir1;
```

**DEFAULT Example** The following statement specifies default auditing options for objects created in the future:

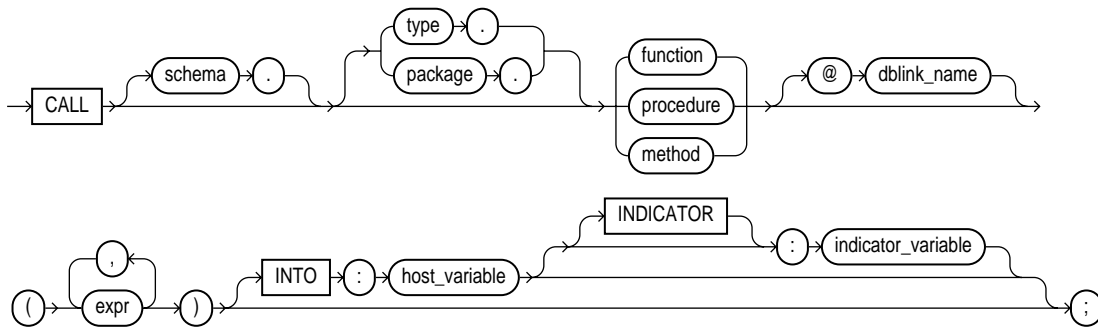
```
AUDIT ALTER, GRANT, INSERT, UPDATE, DELETE
  ON DEFAULT;
```

Any objects created later are automatically audited with the specified options that apply to them, provided that auditing has been enabled:

- If you create a table, Oracle automatically audits any ALTER, GRANT, INSERT, UPDATE, or DELETE statements issued against the table.
- If you create a view, Oracle automatically audits any GRANT, INSERT, UPDATE, or DELETE statements issued against the view.
- If you create a sequence, Oracle automatically audits any ALTER or GRANT statements issued against the sequence.
- If you create a procedure, package, or function, Oracle automatically audits any ALTER or GRANT statements issued against it.

# CALL

## Syntax



## Purpose

Enables you to execute a **routine** (a standalone procedure or function, or a procedure or function defined within a type or package) from within SQL. For information on creating such routine, refer to *PL/SQL User's Guide and Reference*.

## Prerequisites

You must have `EXECUTE` privilege on the standalone routine or on the type or package in which the routine is defined.

## Keywords and Parameters

<i>schema</i>	specifies the schema in which the standalone routine (or the package or type containing the routine) resides. If you do not specify <i>schema</i> , Oracle assumes the routine is in your own schema.
<i>type or package</i>	specifies the type or package in which the routine is defined.
<i>function</i> / <i>procedure</i> / <i>method</i>	specifies the name of the function or procedure being called, or a synonym that translates to a function or procedure.  When you call a type's member function or procedure, if the first argument ( <i>SELF</i> ) is a null <code>IN OUT</code> argument, Oracle returns an error. If <i>SELF</i> is a null <code>IN</code> argument, Oracle returns null. In both cases, the function or procedure is not invoked.  <b>Restriction:</b> If the routine is a function, the <code>INTO</code> clause is mandatory.

---

<code>@dblink</code>	in a distributed database system, specifies the name of the database containing the standalone routine (or the package or functioning containing the routine). If you omit <i>dblink</i> , Oracle looks in your local database.
<code>expr</code>	specifies one or more arguments to the routine. <b>Restrictions:</b> <ul style="list-style-type: none"><li>■ An <i>expr</i> cannot be a pseudocolumn or either of the object reference functions VALUE or REF.</li><li>■ Any <i>expr</i> that is an IN OUT or OUT argument of the routine must correspond to a host variable expression.</li></ul>
<code>INTO :host_variable</code>	applies only to calls to functions. This parameter specifies which host variable will store the return value of the function.
<code>:indicator_variable</code>	indicates the value or condition of the host variable. <b>See Also:</b> <i>Pro*C/C++ Precompiler Programmer's Guide</i> for more information on host variables and indicator variables.

---

## Example

The following statement creates a procedure `UPDATESALARY`, and then calls the procedure, which updates the specified employee `ID` with a new salary.

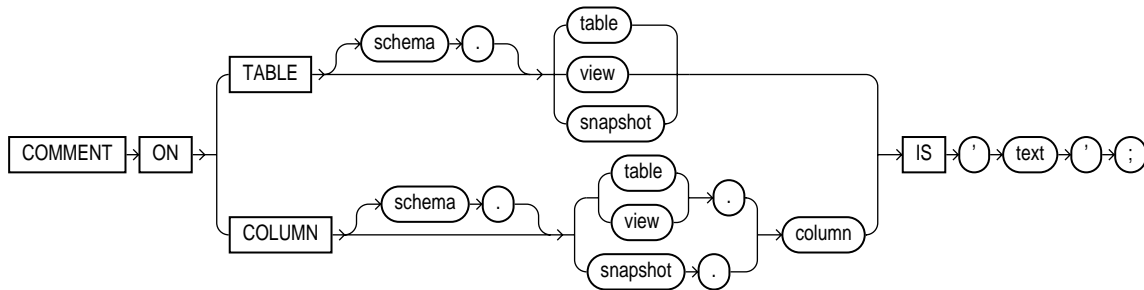
```
CREATE OR REPLACE PROCEDURE updateSalary
  (id NUMBER, newsalary NUMBER) IS
  BEGIN
    UPDATE emp SET sal=newsalary WHERE empno=id;
  END;

CALL updateSalary(1404, 50000);
```

---

## COMMENT

### Syntax



### Purpose

To add a comment about a table, view, snapshot, or column into the data dictionary. See also "[Comments](#)" on page 2-62.

You can view the comments on a particular table or column by querying the data dictionary views `USER_TAB_COMMENTS`, `DBA_TAB_COMMENTS`, or `ALL_TAB_COMMENTS` or `USER_COL_COMMENTS`, `DBA_COL_COMMENTS`, or `ALL_COL_COMMENTS`. For information on these views, see *Oracle8i Reference*.

To drop a comment from the database, set it to the empty string `''`.

### Prerequisites

The table, view, or snapshot must be in your own schema or you must have `COMMENT ANY TABLE` system privilege.

### Keywords and Parameters

---

TABLE	specifies the schema and name of the table, view, or materialized view to be commented. If you omit <i>schema</i> , Oracle assumes the table, view, or snapshot is in your own schema.
COLUMN	specifies the name of the column of a table, view, or snapshot to be commented. If you omit <i>schema</i> , Oracle assumes the table, view, or snapshot is in your own schema.
IS 'text'	is the text of the comment. See the syntax description of 'text' in " <a href="#">Text</a> " on page 2-2.

---



**Example**

To insert an explanatory remark on the `NOTES` column of the `SHIPPING` table, you might issue the following statement:

```
COMMENT ON COLUMN shipping.notes  
    IS 'Special packing or shipping instructions';
```


To drop this comment from the database, issue the following statement:

```
COMMENT ON COLUMN shipping.notes IS ' ';
```

---

## COMMIT

### Syntax



### Purpose

To end your current **transaction** and make permanent all changes performed in the transaction. A transaction is a sequence of SQL statements that Oracle treats as a single unit. This statement also erases all savepoints in the transaction and releases the transaction's locks.

---

---

**Note:** Oracle issues an implicit `COMMIT` before and after any data definition language (DDL) statement.

---

---

You can also use this statement to

- Commit an in-doubt distributed transaction manually
- Terminate a read-only transaction begun by a `SET TRANSACTION` statement.

Oracle Corporation recommends that you explicitly end every transaction in your application programs with a `COMMIT` or `ROLLBACK` statement, including the last transaction, before disconnecting from Oracle. If you do not explicitly commit the transaction and the program terminates abnormally, the last uncommitted transaction is automatically rolled back.

A normal exit from most Oracle utilities and tools causes the current transaction to be committed. A normal exit from an Oracle precompiler program does not commit the transaction and relies on Oracle to roll back the current transaction.

**See Also:**

- *Oracle8i Concepts* for more information on transactions.
- "[SET TRANSACTION](#)" on page 7-602 for more information on specifying characteristics of a transaction.

**Prerequisites**

You need no privileges to commit your current transaction.

To manually commit a distributed in-doubt transaction that you originally committed, you must have `FORCE TRANSACTION` system privilege. To manually commit a distributed in-doubt transaction that was originally committed by another user, you must have `FORCE ANY TRANSACTION` system privilege.

**Keywords and Parameters**


---

WORK	is supported for compliance with standard SQL. The statements <code>COMMIT</code> and <code>COMMIT WORK</code> are equivalent.
COMMENT	<p>specifies a comment to be associated with the current transaction. The 'text' is a quoted literal of up to 50 characters that Oracle stores in the data dictionary view <code>DBA_2PC_PENDING</code> along with the transaction ID if the transaction becomes in-doubt.</p> <p><b>See Also:</b> "<a href="#">COMMENT</a>" on page 7-228 for more information on adding comments to SQL statements.</p>
FORCE	<p>in a distributed database system, manually commits an in-doubt distributed transaction. The transaction is identified by the 'text' containing its local or global transaction ID. To find the IDs of such transactions, query the data dictionary view <code>DBA_2PC_PENDING</code>. You can use <i>integer</i> to specifically assign the transaction a system change number (SCN). If you omit <i>integer</i>, the transaction is committed using the current SCN.</p> <hr/> <p><b>Note:</b> A <code>COMMIT</code> statement with a <code>FORCE</code> clause commits only the specified transaction. Such a statement does not affect your current transaction.</p> <hr/> <p><b>Restriction:</b> <code>COMMIT</code> statements using the <code>FORCE</code> clause are not supported in PL/SQL.</p> <p><b>See Also:</b> <i>Oracle8i Distributed Database Systems</i> for more information on these topics.</p> <hr/>

**Examples**

**INSERT Example** This statement inserts a row into the `DEPT` table and commits this change:

```
INSERT INTO dept VALUES (50, 'MARKETING', 'TAMPA');
```

```
COMMIT WORK;
```

**COMMENT Example** The following statement commits the current transaction and associates a comment with it:

```
COMMIT  
  COMMENT 'In-doubt transaction Code 36, Call (415) 555-2637';
```

If a network or machine failure prevents this distributed transaction from committing properly, Oracle stores the comment in the data dictionary along with the transaction ID. The comment indicates the part of the application in which the failure occurred and provides information for contacting the administrator of the database where the transaction was committed.

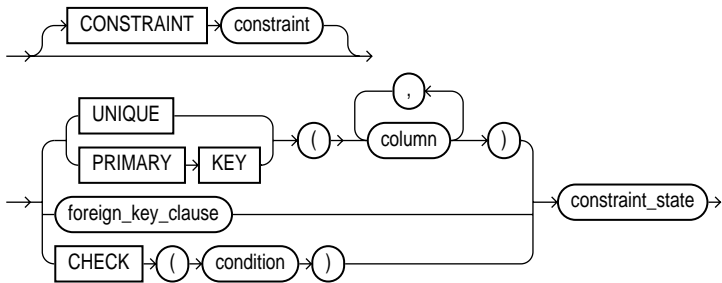
**In-Doubt Transaction Example** The following statement manually commits an in-doubt distributed transaction:

```
COMMIT FORCE '22.57.53';
```

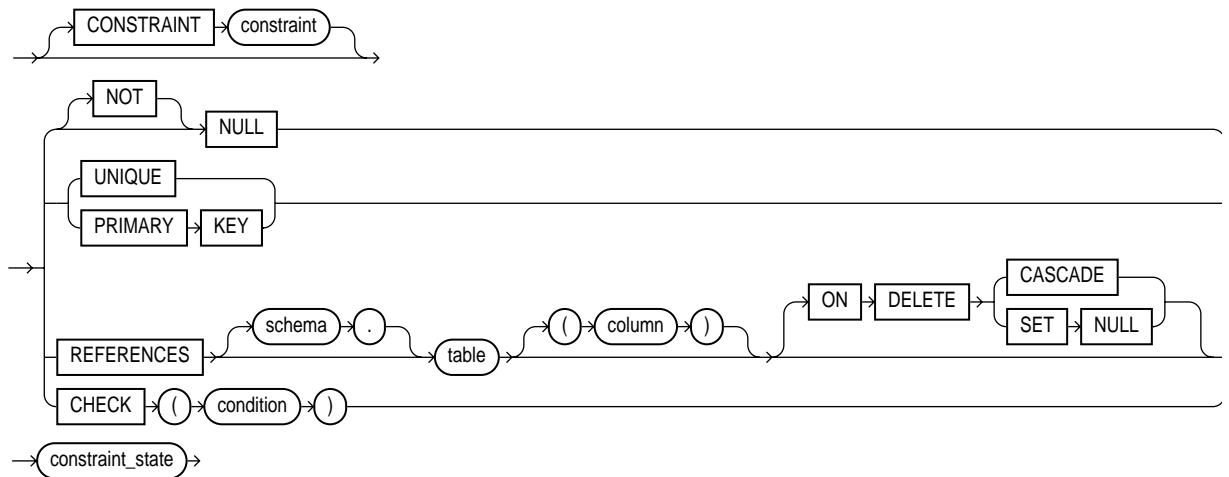
## constraint\_clause

### Syntax

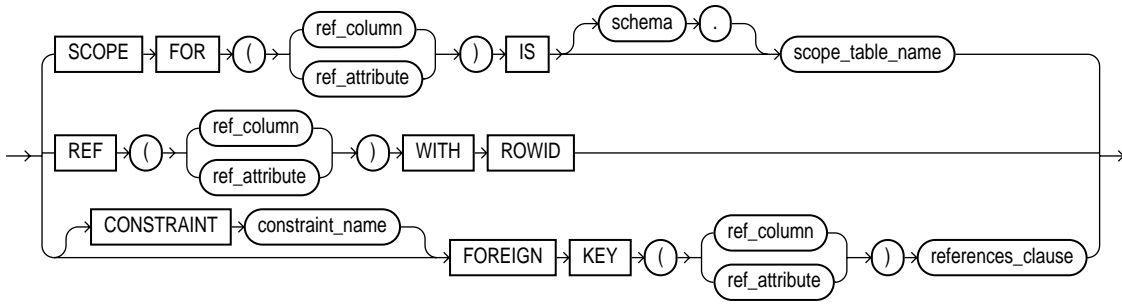
**table\_constraint::=**




**column\_constraint::=**




**table\_ref\_constraint::=**



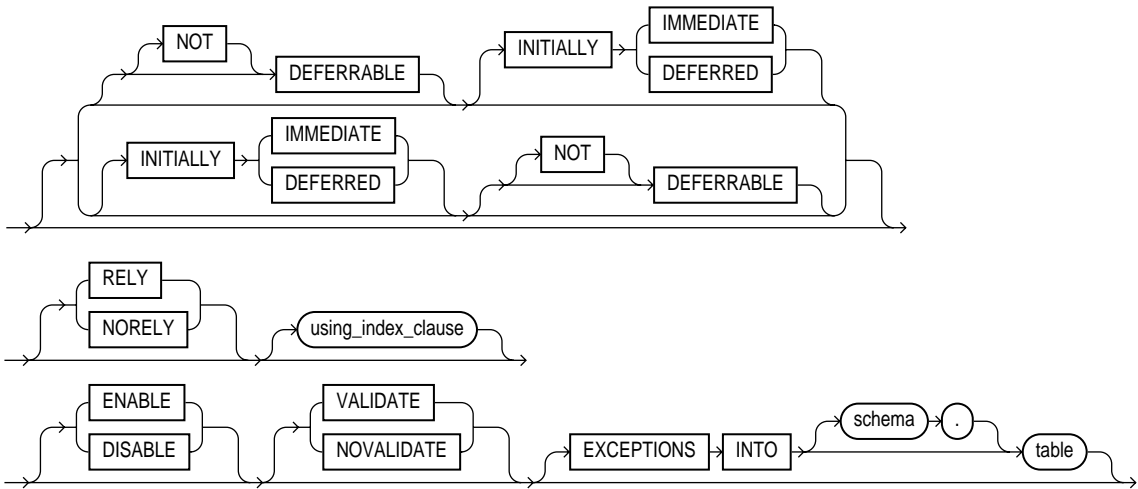
**column\_ref\_constraint::=**



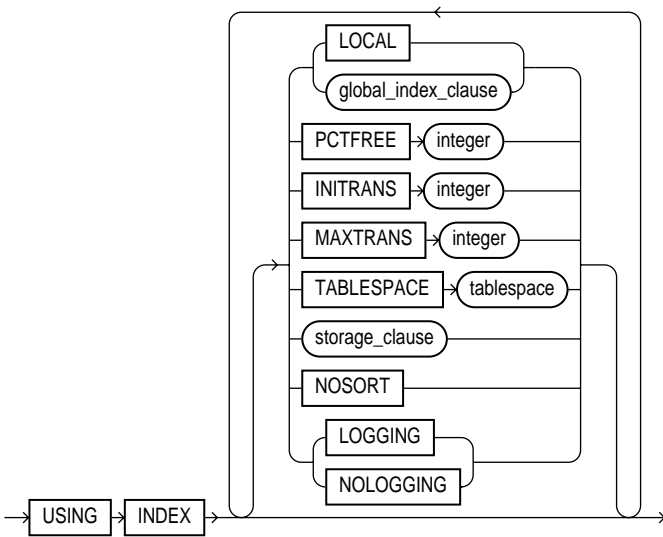
**references\_clause::=**




**constraint\_state::=**



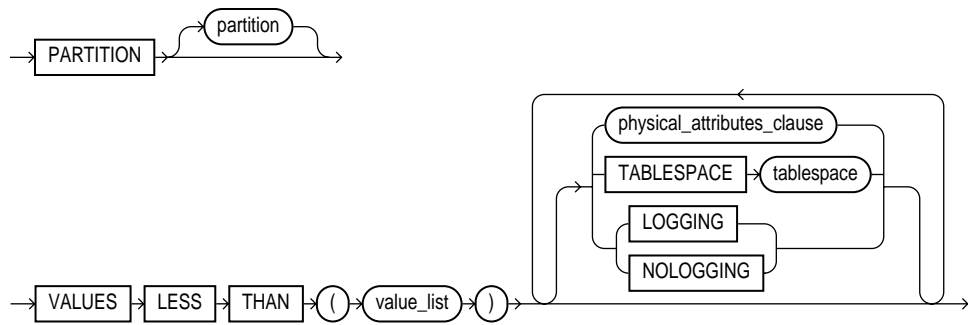
**using\_index\_clause::=**




**global\_index\_clause::=**




**global\_partition\_clause::=**



**foreign\_key\_clause::=**



**physical\_attributes\_clause::=**





**storage\_clause:** See the "storage\_clause" on page 7-605.

## Purpose

To define an integrity constraint. An **integrity constraint** is a rule that restricts the values for one or more columns in a table or an index-organized table.

---



---

**Note:** Oracle does not support constraints on columns or attributes whose type is an object, nested table, varray, REF, or LOB. The only exception is that NOT NULL constraints are supported for columns or attributes whose type is object, VARRAY, REF, or LOB.

---



---

## Prerequisites

Constraint clauses can appear in either CREATE TABLE or ALTER TABLE statements. To define an integrity constraint, you must have the privileges necessary to issue one of these statements. See "CREATE TABLE" on page 4-381 and "ALTER TABLE" on page 7-123.

To create a referential integrity constraint, the parent table must be in your own schema, or you must have the REFERENCES privilege on the columns of the referenced key in the parent table.

## Keywords and Parameters

---

<i>table_constraint</i>	<p>The <i>table_constraint</i> syntax is part of the table definition. An integrity constraint defined with this syntax can impose rules on any columns in the table.</p> <p>The <i>table_constraint</i> syntax can appear in a CREATE TABLE or ALTER TABLE statement. This syntax can define any type of integrity constraint except a NOT NULL constraint.</p>
<i>column_constraint</i>	<p>The <i>column_constraint</i> syntax is part of a column definition. Usually, an integrity constraint defined with this syntax can impose rules only on the column in which it is defined.</p> <ul style="list-style-type: none"> <li>■ The <i>column_constraint</i> syntax that appears in a CREATE TABLE or ALTER TABLE ADD statement can define any type of integrity constraint.</li> <li>■ <i>Column_constraint</i> syntax that appears in an ALTER TABLE MODIFY <i>column_options</i> statement can only define or remove a NOT NULL constraint.</li> </ul>
CONSTRAINT	<p>identifies the integrity constraint by the name <i>constraint</i>. Oracle stores this name in the data dictionary along with the definition of the integrity constraint. If you omit this identifier, Oracle generates a name with the form SYS_C<i>n</i>.</p> <p>If you do not specify NULL or NOT NULL in a column definition, NULL is the default.</p>

**Restriction:** You cannot create a constraint on columns or attributes whose type is user-defined object, LOB, or REF, with the following exceptions:

- You can specify a NOT NULL constraint on columns or attributes of user-defined object type, varray, and LOB.
- You can specify NOT NULL **and** referential integrity constraints on a column of type REF.

## UNIQUE

designates a column or combination of columns as a unique key. To satisfy a UNIQUE constraint, no two rows in the table can have the same value for the unique key. However, the unique key made up of a single column can contain nulls.

A **composite unique key** is made up of a combination of columns. To define a composite unique key, you must use *table\_constraint* syntax rather than *column\_constraint* syntax. Any row that contains nulls in all key columns automatically satisfies the constraint. However, two rows that contain nulls for one or more key columns and the same combination of values for the other key columns violate the constraint.

### Restrictions:

- For a composite unique key, no two rows in the table can have the same combination of values in the key columns.
- A composite unique key cannot have more than 32 columns. The overall size of the key (in bytes) should not exceed approximately the width of all indexed columns plus the number of indexed columns.
- A unique key column cannot be of datatype LONG or LONG RAW.
- You cannot designate the same column or combination of columns as both a unique key and a primary key.

---

**PRIMARY KEY** designates a column or combination of columns as the table's primary key. A **composite primary key** is made up of a combination of columns. To define a composite primary key, you must use the *table\_constraint* syntax rather than the *column\_constraint* syntax.

**Restrictions:**

- A table can have only one primary key.
- None of the columns in the primary key can have datatype LONG, LONG RAW, VARRAY, NESTED TABLE, OBJECT, LOB, BFILE, or REF.
- No primary key value can appear in more than one row in the table.
- No column that is part of the primary key can contain a null.
- The size of the PRIMARY KEY of an index-organized table cannot exceed one-half of the database block size or 3800 bytes, whichever is less. (PRIMARY KEY is required for an index-organized table.)
- A composite primary key cannot have more than 32 columns. The overall size of the key (in bytes) should not exceed approximately the width of all indexed columns plus the number of indexed columns.
- You cannot designate the same column or combination of columns as both a primary key and a unique key.

**NULL | NOT NULL** determines whether a column can contain nulls. You must specify NULL and NOT NULL with *column\_constraint* syntax, not with *table\_constraint* syntax.

**NULL** specifies that a column can contain null values. The NULL keyword does not actually define an integrity constraint. If you do not specify either NOT NULL or NULL, the column can contain nulls by default.

**NOT NULL** specifies that a column cannot contain null values. To satisfy this constraint, every row in the table must contain a value for the column.

**Restriction:** You cannot specify NULL or NOT NULL for an attribute of an object. Instead, use a CHECK constraint with the IS [NOT] NULL condition. See the "[Attribute-Level Constraints Example](#)" on page 7-248.

### Referential integrity constraints

Referential integrity constraints designate a column or combination of columns as the foreign key and establish a relationship between that foreign key and a specified primary or unique key, called the **referenced key**. The table containing the foreign key is called the **child table**, and the table containing the referenced key is called the **parent table**. The foreign key and the referenced key can be in the same table. In this case, the parent and child tables are the same.

From the table level, specify referential integrity using the *foreign\_key\_clause* with the *table\_constraint* syntax. This syntax allows you to specify a **composite foreign key**, which is made up of a combination of columns.

From the column level, use the REFERENCES clause of the *column\_constraint* syntax to specify a referential integrity constraint in which the foreign key is made up of a single column.

You can designate the same column or combination of columns as both a foreign key and a primary or unique key. You can also designate the same column or combination of columns as both a foreign key and a cluster key.

You can define multiple foreign keys in a table. Also, a single column can be part of more than one foreign key.

**Restrictions:**

- A foreign key cannot be of type LONG or LONG RAW.
- The referenced UNIQUE or PRIMARY KEY constraint on the parent table must already be defined.
- The child and parent tables must be on the same database. To enable referential integrity constraints across nodes of a distributed database, you must use database triggers.
- You cannot define a referential integrity constraint in a CREATE TABLE statement that contains an AS *subquery* clause. Instead, you must create the table without the constraint and then add it later with an ALTER TABLE statement.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals*.

*foreign\_key\_clause* designates a column or combination of columns as the foreign key from the table level. You must use this syntax to define a composite foreign key.

To satisfy a referential integrity constraint involving composite keys, either the values of the foreign key columns must match the values of the referenced key columns in a row in the parent table, or the value of at least one of the columns of the foreign key must be null.

**Restrictions:**

- A composite foreign key cannot have more than 32 columns. The overall size of the key (in bytes) should not exceed approximately the width of all indexed columns plus the number of indexed columns.
- A composite foreign key must refer to a composite unique key or a composite primary key.

---

REFERENCES	designates the current column or attribute as the foreign key and identifies the parent table and the column or combination of columns that make up the referenced key. If you identify only the parent table and omit the column names, the foreign key automatically references the primary key of the parent table. The corresponding columns of the referenced key and the foreign key must match in number and datatypes.
ON DELETE	determines how Oracle automatically maintains referential integrity if you remove a referenced primary or unique key value. If you omit this clause, Oracle does not allow you to delete referenced key values in the parent table that have dependent rows in the child table. <ul style="list-style-type: none"> <li>■ CASCADE specifies that Oracle removes dependent foreign key values.</li> <li>■ SET NULL specifies that Oracle converts dependent foreign key values to NULL.</li> </ul>
CHECK	<p>specifies a condition that each row in the table must satisfy. To satisfy the constraint, each row in the table must make the condition either TRUE or unknown (due to a null). For information and syntax, see "<a href="#">Conditions</a>" on page 5-14. When Oracle evaluates a CHECK constraint condition for a particular row, any column names in the condition refer to the column values in that row.</p> <p>If you create multiple CHECK constraints for a column, design them carefully so their purposes do not conflict, and do not assume any particular order of evaluation of the conditions. Oracle does not verify that CHECK conditions are not mutually exclusive.</p> <p><b>Restrictions:</b></p> <p>The condition of a CHECK constraint can refer to any column in the table, but it cannot refer to columns of other tables.</p> <p>CHECK constraint conditions cannot contain the following constructs:</p> <ul style="list-style-type: none"> <li>■ Queries to refer to values in other rows</li> <li>■ Calls to the functions SYSDATE, UID, USER, or USERENV</li> <li>■ The pseudocolumns CURRVAL, NEXTVAL, LEVEL, or ROWNUM</li> <li>■ Date constants that are not fully specified</li> </ul>

*table\_ref\_constraint* and *column\_ref\_constraint* further describe a column of type REF. The only difference between these clauses is that you specify *table\_ref\_constraint* from the table level, so you must identify the REF column or attribute you are defining. You specify *column\_ref\_constraint* after you have already identified the REF column or attribute. Both types of constraint let you specify a SCOPE constraint, a WITH ROWID constraint, or a referential integrity constraint.

As is the case for regular table and column constraints, you use FOREIGN KEY syntax for a referential integrity constraint at the table level, and REFERENCES syntax for a referential integrity constraint at the column level. See "[Referential integrity constraints](#)" on page 7-239.

If the REF column's scope table or reference table has a primary-key-based object identifier, then it is a **user-defined REF column**.

**See Also:** *Oracle8i Concepts* for more information on REFS.

*ref\_column* is the name of a REF column of an object or relational table.

*ref\_attribute* is an embedded REF attribute within an object column of a relational table.

SCOPE In a table with a REF column, each REF value in the column can conceivably reference a row in a different object table. The SCOPE clause restricts the scope of references to a single table, *scope\_table\_name*. The values in the REF column or attribute point to objects in *scope\_table\_name*, in which object instances (of the same type as the REF column) are stored. You can only specify one scope table per REF column.

**Restrictions:**

- You can add a SCOPE constraint to an existing column only if the table is empty.
- You cannot specify SCOPE for the REF elements of a varray column.
- You must specify this clause if you specify AS *subquery* and the subquery returns user-defined REFS.
- The *scope\_table\_name* must be in your own schema or you must have SELECT privileges on *scope\_table\_name* or SELECT ANY TABLE system privileges.
- You cannot drop a SCOPE table constraint from a REF column.

WITH ROWID stores the rowid along with the REF value in *ref\_column* or *ref\_attribute*. Storing a REF value with a rowid can improve the performance of dereferencing operations, but will also use more space. Default storage of REF values is without rowids.

**Restrictions:**

- You cannot specify a `WITH ROWID` constraint for the `REF` elements of a varray column.
- You cannot drop a `WITH ROWID` constraint from a `REF` column.
- If the `REF` column or attribute is scoped, then this clause is ignored and the rowid is not stored with the `REF` value.

*references\_clause* specifies a referential integrity constraint on the `REF` column. This clause also implicitly restricts the scope of the `REF` column or attribute to the reference table.

If you do not specify `CONSTRAINT`, Oracle generates a system name for the constraint.

**Restrictions:**

- If you add a referential integrity constraint to an existing `REF` column that is scoped, then the referenced table must be the same as the scope table of the `REF` column.
- The system adds a scope constraint when you add a referential integrity constraint to an existing unscoped `REF` column. Therefore, all the restrictions that apply for the `SCOPE` constraint also apply in this case.
- If you later drop the referential integrity constraint, the `REF` column will remain scoped to the referenced table.

`DEFERRABLE` indicates that constraint checking can be deferred until the end of the transaction by using the `SET CONSTRAINT(S)` statement. For information on checking constraints after each DML statement, see "[SET CONSTRAINT\(S\)](#)" on page 7-598. See *Oracle8i Administrator's Guide* and *Oracle8i Concepts* for more information about deferred constraints.

`NOT DEFERRABLE` indicates that this constraint is checked at the end of each DML statement. If you do not specify either word, then `NOT DEFERRABLE` is the default.

`INITIALLY IMMEDIATE` indicates that at the start of every transaction, the default is to check this constraint at the end of every DML statement. If you do not specify `INITIALLY`, `INITIALLY IMMEDIATE` is the default.

`INITIALLY DEFERRED` implies that this constraint is `DEFERRABLE` and specifies that, by default, the constraint is checked only at the end of each transaction.

**Restrictions:**

- You cannot defer a `NOT DEFERRABLE` constraint with the `SET CONSTRAINT(S)` statement.
- You cannot specify either `DEFERRABLE` or `NOT DEFERRABLE` if you are modifying an existing constraint directly (that is, by specifying the `ALTER TABLE . . . MODIFY constraint` statement).
- You cannot alter a constraint's deferrability status. You must drop the constraint and re-create it.

**RELY** | **NORELY** specifies whether an enabled constraint is to be enforced. Specify **RELY** to enable an existing constraint without enforcement. Specify **NORELY** to enable and enforce an existing constraint. The default is **NORELY**.

Unenforced constraints are generally useful only with materialized views and query rewrite. Depending on the `QUERY_REWRITE_INTEGRITY` mode (see ["ALTER SESSION"](#) on page 7-83), query rewrite can use constraints that are enabled with or without enforcement to determine join information.

**See Also:** *Oracle8i Data Warehousing Guide* for more information on materialized views and query rewrite.

**Restrictions:**

- **RELY** and **NORELY** are relevant only if you are modifying an existing constraint (that is, you have issued the `ALTER TABLE ... MODIFY constraint` statement).
- You cannot set a `NOT NULL` constraint to **RELY**.

*using\_index\_clause* specifies parameters for the index Oracle uses to enable a `UNIQUE` or `PRIMARY KEY` constraint. The name of the index is the same as the name of the constraint.

You can choose the values of the `INITTRANS`, `MAXTRANS`, `TABLESPACE`, `STORAGE`, and `PCTFREE` parameters for the index. For information on these parameters, see ["CREATE TABLE"](#) on page 4-381.

If *table* is partitioned, you can specify a locally or globally partitioned index for the unique or primary key constraint. For a description of `LOCAL` and *global\_index\_clause*, and for a description of `NOSORT` and `LOGGING` | `NOLOGGING` in relation to indexes, see ["CREATE INDEX"](#) on page 7-291.

**Restriction:** Use this clause only when enabling `UNIQUE` and `PRIMARY KEY` constraints.

**NOSORT** indicates that the rows are stored in the database in ascending order and therefore Oracle does not have to sort the rows when creating the index.

**ENABLE** specifies that the constraint will be applied to all new data in the table. Before you can enable a referential integrity constraint, its referenced constraint must be enabled.

- `ENABLE VALIDATE` additionally specifies that all old data also complies with the constraint. An enabled validated constraint guarantees that all data is and will continue to be valid.

If you place a primary key constraint in `ENABLE VALIDATE` mode, the validation process will verify that the primary key columns contain no nulls. To avoid this overhead, mark each column in the primary key `NOT NULL` before enabling the table's primary key constraint. (For optimal results, do this before inserting data into the column.)

- `ENABLE NOVALIDATE` ensures that all new DML operations on the constrained data comply with the constraint, but does not ensure that existing data in the table complies with the constraint.



Enabling a primary key or unique key constraint automatically creates a unique index to enforce the constraint. This index is dropped if the constraint is subsequently disabled, causing Oracle to rebuild the index every time the constraint is enabled. To avoid this behavior, create new primary key and unique key constraints initially disabled. Then create nonunique indexes or use existing nonunique indexes to enforce the constraints.

For additional notes and restrictions, see the [enable\\_disable\\_clause](#) of "CREATE TABLE" on page 4-405.

## DISABLE

disables the integrity constraint. If you do not specify this clause when creating a constraint, Oracle automatically enables the constraint.

- `DISABLE VALIDATE` disables the constraint and drops the index on the constraint, but keeps the constraint valid. This feature is most useful in data warehousing situations, where the need arises to load into a range-partitioned table a quantity of data with a distinct range of values in the unique key. In such situations, the disable validate state enables you to save space by not having an index. You can then load data from a non-partitioned table into a partitioned table using the [exchange\\_partition\\_clause](#) of the `ALTER TABLE` statement or using SQL\*Loader. **All other modifications** to the table (inserts, updates, and deletes) by other SQL statements are disallowed.

If the unique key coincides with the partitioning key of the partitioned table, disabling the constraint saves overhead and has no detrimental effects. If the unique key does not coincide with the partitioning key, Oracle performs automatic table scans during the exchange to validate the constraint, which might offset the benefit of loading without an index.

- `DISABLE NOVALIDATE` signifies that Oracle makes no effort to maintain the constraint (because it is disabled) and cannot guarantee that the constraint is true (because it is not being validated). For information on when to use this setting, see *Oracle8i Designing and Tuning for Performance*.

You cannot drop a table whose primary key is being referenced by a foreign key even if the foreign key constraint is in `DISABLE NOVALIDATE` state. Further, the optimizer can use constraints in `DISABLE NOVALIDATE` state.

- If you specify neither `VALIDATE` nor `NOVALIDATE`, the default is `NOVALIDATE`.
- If you disable a unique or primary key constraint that is using a unique index, Oracle drops the unique index.

EXCEPTIONS  
INTO

specifies a table into which Oracle places the rowids of all rows violating the constraint. If you omit schema, Oracle assumes the exceptions table is in your own schema. If you omit this clause altogether, Oracle assumes that the table is named `EXCEPTIONS`. The exceptions table must be on your local database.

---

You can create the `EXCEPTIONS` table using one of these scripts:

- `UTLEXCPT.SQL` uses physical rowids. Therefore it can accommodate rows from conventional tables but not from index-organized tables. (See the Note that follows.)
- `UTLEXP1.SQL` uses universal rowids, so it can accommodate rows from both conventional and index-organized tables.

If you create your own exceptions table, it must follow the format prescribed by one of these two scripts. See *Oracle8i Migration* for compatibility issues related to the use of these scripts.

---

**Note:** If you are collecting exceptions from index-organized tables based on primary keys (rather than universal rowids), you must create a separate exceptions table for each index-organized table to accommodate its primary-key storage. You create multiple exceptions tables with different names by modifying and resubmitting the script.

For information on the SQL scripts, see the `DBMS_IOT` package in *Oracle8i Supplied PL/SQL Packages Reference*. For information on eliminating migrated and chained rows, see *Oracle8i Designing and Tuning for Performance*.

---

**This clause is valid only when validating a constraint.**

---

## Examples

**Unique Key Example** The following statement creates the `DEPT` table and defines and enables a unique key on the `DNAME` column:

```
CREATE TABLE dept
  (deptno NUMBER(2),
   dname  VARCHAR2(9) CONSTRAINT unq_dname UNIQUE,
   loc    VARCHAR2(10) );
```

The constraint `UNQ_DNAME` identifies the `DNAME` column as a unique key. This constraint ensures that no two departments in the table have the same name. However, the constraint does allow departments without names.

Alternatively, you can define and enable this constraint with the *table\_constraint* syntax:

```
CREATE TABLE dept
  (deptno NUMBER(2),
   dname  VARCHAR2(9),
   loc    VARCHAR2(10),
   CONSTRAINT unq_dname
   UNIQUE (dname)
 USING INDEX PCTFREE 20
   TABLESPACE user_x
```

```
STORAGE ( INITIAL 8K NEXT 6K ) );
```

The above statement also uses the `USING INDEX` clause to specify storage characteristics for the index that Oracle creates to enable the constraint.

**Composite Unique Key Example** The following statement defines and enables a composite unique key on the combination of the `CITY` and `STATE` columns of the `CENSUS` table:

```
ALTER TABLE census
  ADD CONSTRAINT unq_city_state
  UNIQUE (city, state)
  USING INDEX PCTFREE 5
  TABLESPACE user_y
  EXCEPTIONS INTO bad_keys_in_ship_cont;
```

The `UNQ_CITY_STATE` constraint ensures that the same combination of `CITY` and `STATE` values does not appear in the table more than once.

The `ADD CONSTRAINT` clause also specifies other properties of the constraint:

- The `USING INDEX` clause specifies storage characteristics for the index Oracle creates to enable the constraint.
- The `EXCEPTIONS INTO` clause causes Oracle to write information to the `BAD_KEYS_IN_SHIP_CONT` table about any rows currently in the `CENSUS` table that violate the constraint.

**Primary Key Example** The following statement creates the `DEPT` table and defines and enables a primary key on the `DEPTNO` column:

```
CREATE TABLE dept
  (deptno NUMBER(2) CONSTRAINT pk_dept PRIMARY KEY,
  dname VARCHAR2(9),
  loc VARCHAR2(10) );
```

The `PK_DEPT` constraint identifies the `DEPTNO` column as the primary key of the `DEPT` table. This constraint ensures that no two departments in the table have the same department number and that no department number is `NULL`.

Alternatively, you can define and enable this constraint with *table\_constraint* syntax:

```
CREATE TABLE dept
  (deptno NUMBER(2),
  dname VARCHAR2(9),
  loc VARCHAR2(10),
```

```
CONSTRAINT pk_dept PRIMARY KEY (deptno) );
```

**Composite Primary Key Example** The following statement defines a composite primary key on the combination of the SHIP\_NO and CONTAINER\_NO columns of the SHIP\_CONT table:

```
ALTER TABLE ship_cont
  ADD PRIMARY KEY (ship_no, container_no) DISABLE;
```

This constraint identifies the combination of the SHIP\_NO and CONTAINER\_NO columns as the primary key of the SHIP\_CONT table. The constraint ensures that no two rows in the table have the same values for both the SHIP\_NO column and the CONTAINER\_NO column.

The CONSTRAINT clause also specifies the following properties of the constraint:

- The constraint definition does not include a constraint name, so Oracle generates a name for the constraint.
- The DISABLE clause causes Oracle to define the constraint but not enable it.

**NOT NULL Example** The following statement alters the EMP table and defines and enables a NOT NULL constraint on the SAL column:

```
ALTER TABLE emp
  MODIFY (sal NUMBER CONSTRAINT nn_sal NOT NULL);
```

NN\_SAL ensures that no employee in the table has a null salary.

**Attribute-Level Constraints Example** The following example guarantees that a value exists for both the FIRST\_NAME and LAST\_NAME attributes of the NAME column in the STUDENTS table:

```
CREATE TYPE person_name AS OBJECT
  (first_name VARCHAR2(30), last_name VARCHAR2(30));

CREATE TABLE students (name person_name, age INTEGER,
  CHECK (name.first_name IS NOT NULL AND
  name.last_name IS NOT NULL));
```

**Referential Integrity Constraint Example** The following statement creates the EMP table and defines and enables a foreign key on the DEPTNO column that references the primary key on the DEPTNO column of the DEPT table:

```
CREATE TABLE emp
  (empno NUMBER(4),
```

```

ename      VARCHAR2(10),
job        VARCHAR2(9),
mgr         NUMBER(4),
hiredate   DATE,
sal        NUMBER(7,2),
comm       NUMBER(7,2),
deptno     CONSTRAINT fk_deptno REFERENCES dept(deptno) );

```

The constraint `FK_DEPTNO` ensures that all departments given for employees in the `EMP` table are present in the `DEPT` table. However, employees can have null department numbers, meaning they are not assigned to any department. To ensure that all employees are assigned to a department, you could create a `NOT NULL` constraint on the `DEPTNO` column in the `EMP` table, in addition to the `REFERENCES` constraint.

Before you define and enable this constraint, you must define and enable a constraint that designates the `DEPTNO` column of the `DEPT` table as a primary or unique key.

The referential integrity constraint definition does not use the `FOREIGN KEY` keyword to identify the columns that make up the foreign key. Because the constraint is defined with a column constraint clause on the `DEPTNO` column, the foreign key is automatically on the `DEPTNO` column.

The constraint definition identifies both the parent table and the columns of the referenced key. Because the referenced key is the parent table's primary key, the referenced key column names are optional.

The above statement omits the `DEPTNO` column's datatype. Because this column is a foreign key, Oracle automatically assigns it the datatype of the `DEPT.DEPTNO` column to which the foreign key refers.

Alternatively, you can define a referential integrity constraint with *table\_constraint* syntax:

```

CREATE TABLE emp
(empno      NUMBER(4),
ename      VARCHAR2(10),
job        VARCHAR2(9),
mgr         NUMBER(4),
hiredate   DATE,
sal        NUMBER(7,2),
comm       NUMBER(7,2),
deptno,
CONSTRAINT fk_deptno
FOREIGN KEY (deptno)

```

```
REFERENCES dept(deptno) );
```

The foreign key definitions in both statements of this statement omit the `ON DELETE` clause, causing Oracle to forbid the deletion of a department if any employee works in that department.

**ON DELETE Example** This statement creates the `EMP` table, defines and enables two referential integrity constraints, and uses the `ON DELETE` clause:

```
CREATE TABLE emp
(empno NUMBER(4) PRIMARY KEY,
ename VARCHAR2(10),
job VARCHAR2(9),
mgr NUMBER(4) CONSTRAINT fk_mgr
REFERENCES emp ON DELETE SET NULL,
hiredate DATE,
sal NUMBER(7,2),
comm NUMBER(7,2),
deptno NUMBER(2) CONSTRAINT fk_deptno
REFERENCES dept(deptno)
ON DELETE CASCADE );
```

Because of the first `ON DELETE` clause, if manager number 2332 is deleted from the `EMP` table, Oracle sets to null the value of `MGR` for all employees in the `EMP` table who previously had manager 2332.

Because of the second `ON DELETE` clause, Oracle cascades any deletion of a `DEPTNO` value in the `DEPT` table to the `DEPTNO` values of its dependent rows of the `EMP` table. For example, if Department 20 is deleted from the `DEPT` table, Oracle deletes the department's employees from the `EMP` table.

**Composite Referential Integrity Constraint Example** The following statement defines and enables a foreign key on the combination of the `AREACO` and `PHONENO` columns of the `PHONE_CALLS` table:

```
ALTER TABLE phone_calls
ADD CONSTRAINT fk_areaco_phoneno
FOREIGN KEY (areaco, phoneno)
REFERENCES customers(areaco, phoneno)
EXCEPTIONS INTO wrong_numbers;
```

The constraint `FK_AREACO_PHONENO` ensures that all the calls in the `PHONE_CALLS` table are made from phone numbers that are listed in the `CUSTOMERS` table. Before you define and enable this constraint, you must define and enable a constraint that

designates the combination of the AREACO and PHONENO columns of the CUSTOMERS table as a primary or unique key.

The EXCEPTIONS INTO clause causes Oracle to write information to the WRONG\_NUMBERS table about any rows in the PHONE\_CALLS table that violate the constraint.

**CHECK Constraint Examples** The following statement creates the DEPT table and defines a CHECK constraint in each of the table's columns:

```
CREATE TABLE dept
  (deptno NUMBER CONSTRAINT check_deptno
    CHECK (deptno BETWEEN 10 AND 99)
    DISABLE,
   dname VARCHAR2(9) CONSTRAINT check_dname
    CHECK (dname = UPPER(dname))
    DISABLE,
   loc VARCHAR2(10) CONSTRAINT check_loc
    CHECK (loc IN ('DALLAS', 'BOSTON',
    'NEW YORK', 'CHICAGO'))
    DISABLE);
```

Each constraint restricts the values of the column in which it is defined:

- CHECK\_DEPTNO ensures that no department numbers are less than 10 or greater than 99.
- CHECK\_DNAME ensures that all department names are in uppercase.
- CHECK\_LOC restricts department locations to Dallas, Boston, New York, or Chicago.

Because each CONSTRAINT clause contains the DISABLE clause, Oracle only defines the constraints and does not enable them.

The following statement creates the EMP table and uses a *table\_constraint\_clause* to define and enable a CHECK constraint:

```
CREATE TABLE emp
  (empno          NUMBER(4),
   ename          VARCHAR2(10),
   job            VARCHAR2(9),
   mgr            NUMBER(4),
   hiredate       DATE,
   sal            NUMBER(7,2),
   comm           NUMBER(7,2),
   deptno         NUMBER(2),
```

```
CHECK (sal + comm <= 5000) );
```

This constraint uses an inequality condition to limit an employee's total compensation, the sum of salary and commission, to \$5000:

- If an employee has non-null values for both salary and commission, the sum of these values must not exceed \$5000 to satisfy the constraint.
- If an employee has a null salary or commission, the result of the condition is unknown and the employee automatically satisfies the constraint.

Because the `CONSTRAINT` clause in this example does not supply a constraint name, Oracle generates a name for the constraint.

The following statement defines and enables a `PRIMARY KEY` constraint, two referential integrity constraints, a `NOT NULL` constraint, and two `CHECK` constraints:

```
CREATE TABLE order_detail
(CONSTRAINT pk_od PRIMARY KEY (order_id, part_no),
 order_id NUMBER
    CONSTRAINT fk_oid REFERENCES scott.order (order_id),
 part_no      NUMBER
    CONSTRAINT fk_pno REFERENCES scott.part (part_no),
 quantity    NUMBER
    CONSTRAINT nn_qty NOT NULL
    CONSTRAINT check_qty_low CHECK (quantity > 0),
 cost        NUMBER
    CONSTRAINT check_cost CHECK (cost > 0) );
```

The constraints enable the following rules on table data:

- `PK_OD` identifies the combination of the `ORDER_ID` and `PART_NO` columns as the primary key of the table. To satisfy this constraint, no two rows in the table can contain the same combination of values in the `ORDER_ID` and the `PART_NO` columns, and no row in the table can have a null in either the `ORDER_ID` column or the `PART_NO` column.
- `FK_OID` identifies the `ORDER_ID` column as a foreign key that references the `ORDER_ID` column in the `ORDER` table in `SCOTT`'s schema. All new values added to the column `ORDER_DETAIL.ORDER_ID` must already appear in the column `SCOTT.ORDER.ORDER_ID`.
- `FK_PNO` identifies the `PART_NO` column as a foreign key that references the `PART_NO` column in the `PART` table owned by `SCOTT`. All new values added to the column `ORDER_DETAIL.PART_NO` must already appear in the column `SCOTT.PART.PART_NO`.



- `NN_QTY` forbids nulls in the `QUANTITY` column.
- `CHECK_QTY` ensures that values in the `QUANTITY` column are always greater than zero.
- `CHECK_COST` ensures the values in the `COST` column are always greater than zero.

This example also illustrates the following points about constraint clauses and column definitions:

- *Table\_constraint* syntax and column definitions can appear in any order. In this example, the *table\_constraint* syntax that defines the `PK_OD` constraint precedes the column definitions.
- A column definition can use *column\_constraint* syntax multiple times. In this example, the definition of the `QUANTITY` column contains the definitions of both the `NN_QTY` and `CHECK_QTY` constraints.
- A table can have multiple `CHECK` constraints. Multiple `CHECK` constraints, each with a simple condition enforcing a single business rule, is better than a single `CHECK` constraint with a complicated condition enforcing multiple business rules. When a constraint is violated, Oracle returns an error identifying the constraint. Such an error more precisely identifies the violated business rule if the identified constraint enables a single business rule.

**DEFERRABLE Constraint Examples** The following statement creates table `GAMES` with a `NOT DEFERRABLE INITIALLY IMMEDIATE` constraint check on the `SCORES` column:

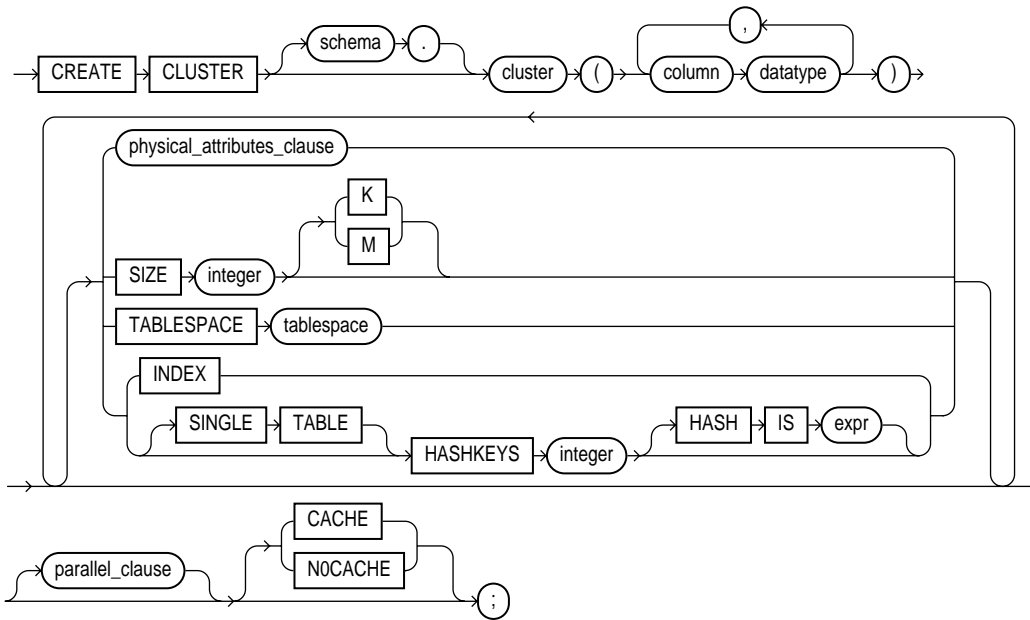
```
CREATE TABLE games (scores NUMBER CHECK (scores >= 0));
```

To define a unique constraint on a column as `INITIALLY DEFERRED DEFERRABLE`, issue the following statement:


```
CREATE TABLE orders
(ord_num NUMBER CONSTRAINT unq_num UNIQUE (ord_num)
INITIALLY DEFERRED DEFERRABLE);
```

# CREATE CLUSTER

## Syntax




### physical\_attributes\_clause::=



**storage\_clause:** See the "[storage\\_clause](#)" on page 7-605.

**parallel\_clause::=**



## Purpose

To create a cluster. A **cluster** is a schema object that contains data from one or more tables, all of which have one or more columns in common. Oracle stores together all the rows (from all the tables) that share the same cluster key.

For general information on clusters, see *Oracle8i Concepts*. For information on performance considerations of clusters, see *Oracle8i Application Developer's Guide - Fundamentals*. For suggestions on when to use clusters, see *Oracle8i Designing and Tuning for Performance*.

## Prerequisites

To create a cluster in your own schema, you must have `CREATE CLUSTER` system privilege. To create a cluster in another user's schema, you must have `CREATE ANY CLUSTER` system privilege. Also, the owner of the schema to contain the cluster must have either space quota on the tablespace containing the cluster or `UNLIMITED TABLESPACE` system privilege.

Oracle does not automatically create an index for a cluster when the cluster is initially created. Data manipulation language (DML) statements cannot be issued against clustered tables until a cluster index has been created.

## Keywords and Parameters

---

<i>schema</i>	is the schema to contain the cluster. If you omit <i>schema</i> , Oracle creates the cluster in your current schema.
<i>cluster</i>	is the name of the cluster to be created.

After you create a cluster, you add tables to it. A cluster can contain a maximum of 32 tables. After you create a cluster and add tables to it, the cluster is transparent. You can access clustered tables with SQL statements just as you can nonclustered tables. For information on adding tables to a cluster, see "[CREATE TABLE](#)" on page 4-381.

<i>column</i>	<p>is the name of a column in the cluster key. You can specify up to 16 cluster key columns. These columns must correspond in both datatype and size to columns in each of the clustered tables, although they need not correspond in name.</p> <p>You cannot specify integrity constraints as part of the definition of a cluster key column. Instead, you can associate integrity constraints with the tables that belong to the cluster.</p>
<i>datatype</i>	<p>is the datatype of a cluster key column. For information on datatypes, see the section "<a href="#">Datatypes</a>" on page 2-9.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot specify a cluster key column of datatype LONG, LONG RAW, REF, nested table, varray, BLOB, CLOB, BFILE, or user-defined object type.</li><li>■ You cannot use the HASH IS clause if any column datatype is not INTEGER or NUMBER with scale 0.</li><li>■ You can specify a column of type ROWID, but Oracle does not guarantee that the values in such columns are valid rowids.</li></ul>
<i>physical_attributes_clause</i>	<p>specifies the storage characteristics of the cluster. Each table in the cluster uses these storage characteristics as well.</p>
PCTUSED	<p>specifies the limit that Oracle uses to determine when additional rows can be added to a cluster's data block. The value of this parameter is expressed as a whole number and interpreted as a percentage.</p>
PCTFREE	<p>specifies the space reserved in each of the cluster's data blocks for future expansion. The value of the parameter is expressed as a whole number and interpreted as a percentage.</p>
INITRANS	<p>specifies the initial number of concurrent update transactions allocated for data blocks of the cluster. The value of this parameter for a cluster cannot be less than 2 or more than the value of the MAXTRANS parameter. The default value is 2 or the INITRANS value for the cluster's tablespace, whichever is greater.</p>
MAXTRANS	<p>specifies the maximum number of concurrent update transactions for any given data block belonging to the cluster. The value of this parameter cannot be less than the value of the INITRANS parameter. The maximum value of this parameter is 255. The default value is the MAXTRANS value for the tablespace to contain the cluster.</p> <p>For a complete description of the PCTUSED, PCTFREE, INITRANS, and MAXTRANS parameters, see "<a href="#">CREATE TABLE</a>" on page 4-381.</p>
<i>storage_clause</i>	<p>specifies how data blocks are allocated to the cluster. See the "<a href="#">storage_clause</a>" on page 7-605.</p>
SIZE	<p>specifies the amount of space in bytes to store all rows with the same cluster key value or the same hash value. Use K or M to specify this space in kilobytes or megabytes. This space determines the maximum number of cluster or hash values stored in a data block. If SIZE is not a divisor of the data block size, Oracle uses the next largest divisor. If SIZE is larger than the data block size, Oracle uses the operating system block size, reserving at least one data block per cluster or hash value.</p>

---

Oracle also considers the length of the cluster key when determining how much space to reserve for the rows having a cluster key value. Larger cluster keys require larger sizes. To see the actual size, query the `KEY_SIZE` column of the `USER_CLUSTERS` data dictionary view. (This does not apply to hash clusters, because hash values are not actually stored in the cluster.)

If you omit this parameter, Oracle reserves one data block for each cluster key value or hash value.

`TABLESPACE` specifies the tablespace in which the cluster is created.

`INDEX` creates an **indexed cluster**. In an indexed cluster, Oracle stores together rows having the same cluster key value. Each distinct cluster key value is stored only once in each data block, regardless of the number of tables and rows in which it occurs.

After you create an indexed cluster, you must create an index on the cluster key before you can issue any data manipulation language (DML) statements against a table in the cluster. This index is called the *cluster index*. For information on creating a cluster index, see "[CREATE INDEX](#)" on page 7-291.

---

**Note:** You cannot create a cluster index for a hash cluster, and you need not create an index on a hash cluster key. If you specify neither `INDEX` nor `HASHKEYS`, Oracle creates an indexed cluster by default.

---

**See Also:** *Oracle8i Concepts* for more information in indexed clusters.

`HASHKEYS` creates a **hash cluster** and specifies the number of hash values for a hash cluster. In a hash cluster, Oracle stores together rows that have the same hash key value. The hash value for a row is the value returned by the cluster's hash function.

Oracle rounds up the `HASHKEYS` value to the nearest prime number to obtain the actual number of hash values. The minimum value for this parameter is 2. If you omit both the `INDEX` clause and the `HASHKEYS` parameter, Oracle creates an indexed cluster by default.

When you create a hash cluster, Oracle immediately allocates space for the cluster based on the values of the `SIZE` and `HASHKEYS` parameters.

**See Also:** *Oracle8i Concepts* for more information on how Oracle allocates space for clusters.

`SINGLE TABLE` specifies that the cluster is a type of hash cluster containing only one table. This clause can provide faster access to rows than would result if the table were not part of a cluster.

**Restriction:** Only one table can be present in the cluster at a time. However, you can drop the table and create a different table in the same cluster.

`HASH IS` specifies an expression to be used as the hash function for the hash cluster. The expression:

- Must evaluate to a positive value
- Must contain at least one column with referenced columns of any datatype as long as the entire expression evaluates to a number of scale 0. For example: `NUM_COLUMN * length(VARCHAR2_COLUMN)`
- Cannot reference user-defined PL/SQL functions
- Cannot reference `SYSDATE`, `USERENV`, `TO_DATE`, `UID`, `USER`, `LEVEL`, or `ROWNUM`
- Cannot evaluate to a constant
- Cannot contain a subquery
- Cannot contain columns qualified with a schema or object name (other than the cluster name)

If you omit the `HASH IS` clause, Oracle uses an internal hash function for the hash cluster.

The cluster key of a hash column can have one or more columns of any datatype. Hash clusters with composite cluster keys or cluster keys made up of noninteger columns must use the internal hash function.

*parallel\_clause* causes creation of the cluster to be parallelized. See also the [Notes](#) to the *parallel\_clause* of "CREATE TABLE" on page 4-381.

`NOPARALLEL` specifies serial execution. This is the default.

`PARALLEL` causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the `PARALLEL_THREADS_PER_CPU` initialization parameter.

`PARALLEL integer` specifies the **degree of parallelism**, which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution servers. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify *integer*.

**Restriction:** If the tables in *cluster* contain any columns of LOB or user-defined object type, this statement as well as subsequent `INSERT`, `UPDATE`, or `DELETE` operations on *cluster* are executed serially without notification.

`CACHE` specifies that the blocks retrieved for this table are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This clause is useful for small lookup tables.

`NOCACHE` specifies that the blocks retrieved for this table are placed at the least recently used end of the LRU list in the buffer cache when a full table scan is performed. This is the default behavior.

---

## Examples

**Creating a Cluster** The following statement creates an indexed cluster named **PERSONNEL** with the cluster key column **DEPARTMENT\_NUMBER**, a cluster size of 512 bytes, and storage parameter values:

```
CREATE CLUSTER personnel
  ( department_number NUMBER(2) )
  SIZE 512
  STORAGE ( INITIAL 100K NEXT 50K );
```

**Adding Tables to a Cluster** The following statements add the **EMP** and **DEPT** tables to the cluster:

```
CREATE TABLE emp
  ( empno      NUMBER           PRIMARY KEY,
    ename      VARCHAR2(10)    NOT NULL
                                CHECK (ename = UPPER(ename)),
    job        VARCHAR2(9),
    mgr        NUMBER           REFERENCES scott.emp(empno),
    hiredate   DATE
              CHECK (hiredate < TO_DATE ('08-14-1998', 'MM-DD-YYYY')),
    sal        NUMBER(10,2)    CHECK (sal > 500),
    comm       NUMBER(9,0)     DEFAULT NULL,
    deptno     NUMBER(2)       NOT NULL )
  CLUSTER personnel (deptno);

CREATE TABLE dept
  ( deptno     NUMBER(2),
    dname      VARCHAR2(9),
    loc        VARCHAR2(9) )
  CLUSTER personnel (deptno);
```

**Cluster Key Example** The following statement creates the cluster index on the cluster key of **PERSONNEL**:

```
CREATE INDEX idx_personnel ON CLUSTER personnel;
```

After creating the cluster index, you can insert rows into either the **EMP** or **DEPT** tables.

**Hash Cluster Examples** The following statement creates a hash cluster named **PERSONNEL** with the cluster key column **DEPARTMENT\_NUMBER**, a maximum of 503 hash key values, each of which is allocated 512 bytes, and storage parameter values:

```
CREATE CLUSTER personnel
( department_number NUMBER )
  SIZE 512 HASHKEYS 500
  STORAGE (INITIAL 100K NEXT 50K);
```

Because the above statement omits the `HASH IS` clause, Oracle uses the internal hash function for the cluster.

The following statement creates a hash cluster named `PERSONNEL` with the cluster key made up of the columns `HOME_AREA_CODE` and `HOME_PREFIX`, and uses a SQL expression containing these columns for the hash function:

```
CREATE CLUSTER personnel
( home_area_code NUMBER,
  home_prefix    NUMBER )
  HASHKEYS 20
  HASH IS MOD(home_area_code + home_prefix, 101);
```


**Single-Table Hash Cluster Example** The following statement creates a single-table hash cluster named `PERSONNEL` with the cluster key `DEPTNO` and a maximum of 503 hash key values, each of which is allocated 512 bytes:

```
CREATE CLUSTER personnel
( deptno NUMBER )
  SIZE 512 SINGLE TABLE HASHKEYS 500;
```



# CREATE CONTEXT

## Syntax



## Purpose

To create a namespace for a **context** (a set of application-defined attributes that validates and secures an application) and to associate the namespace with the externally created package that sets the context.

**See Also:** *Oracle8i Concepts* for a definition and discussion of contexts.

## Prerequisites

To create a context namespace, you must have `CREATE ANY CONTEXT` system privilege.

## Keywords and Parameters

<code>OR REPLACE</code>	redefines an existing context namespace using a different package.
<i>namespace</i>	is the name of the context namespace to create or modify. Context namespaces are always stored in the schema <code>SYS</code> .
<i>schema</i>	is the schema owning <i>package</i> . If you omit <i>schema</i> , Oracle uses the current schema.
<i>package</i>	is the PL/SQL package that sets or resets the context attributes under the namespace for a user session. <b>See Also:</b> <i>Oracle8i Supplied PL/SQL Packages Reference</i> for more information on setting the package.

**Note:** To provide some design flexibility, Oracle does not verify the existence of the schema or the validity of the package at the time you create the context.

## Examples

Suppose you have a human resources (HR) application and a PL/SQL package (HR\_SECURE\_CONTEXT), which validates and secures the HR application. The following statement creates the context namespace HR\_CONTEXT and associates it with the package HR\_SECURE\_CONTEXT:

```
CREATE CONTEXT hr_context USING hr_secure_context;
```

You can control data access based on this context using the SYS\_CONTEXT function. For example, suppose your HR\_SECURE\_CONTEXT package has defined an attribute ORG\_ID as a particular organization identifier. You can secure a base table HR\_ORG\_UNIT by creating a view that restricts access based on the value of ORG\_ID, as follows:

```
CREATE VIEW hr_org_secure_view AS
  SELECT * FROM hr_org_unit
  WHERE organization_id = SYS_CONTEXT('hr_context', 'org_id');
```

**See Also:** ["SYS\\_CONTEXT"](#) on page 4-93 for more information on the SYS\_CONTEXT function.

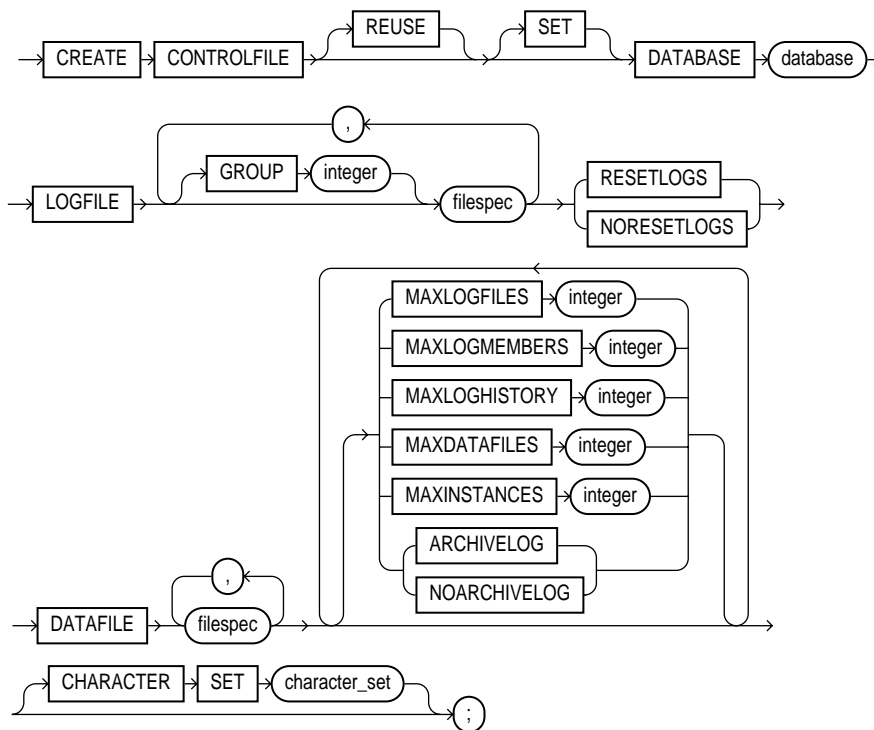
## CREATE CONTROLFILE

### Syntax

---

**WARNING:** Oracle recommends that you perform a full backup of all files in the database before using this statement. For more information, see *Oracle8i Backup and Recovery Guide*.

---



**filespec:** See "filespec" on page 7-516.

### Purpose

To re-create a control file in one of the following cases:

- All copies of your existing control files have been lost through media failure.
- You want to change the name of the database.
- You want to change the maximum number of redo log file groups, redo log file members, archived redo log files, datafiles, or instances that can concurrently have the database mounted and open.

When you issue a `CREATE CONTROLFILE` statement, Oracle creates a new control file based on the information you specify in the statement. If you omit any clauses, Oracle uses the default values rather than the values for the previous control file. After successfully creating the control file, Oracle mounts the database in the mode specified by the initialization parameter `PARALLEL_SERVER`. You then must perform media recovery before opening the database. It is recommended that you then shut down the instance and take a full backup of all files in the database.

**See Also:** *Oracle8i Backup and Recovery Guide* for more information about using this statement.

### Prerequisites

You must have the `OSDBA` role enabled. The database must not be mounted by any instance.

If the `REMOTE_LOGIN_PASSWORDFILE` initialization parameter is set to `EXCLUSIVE`, Oracle returns an error when you attempt to re-create the control file. To avoid this message, either set the parameter to `SHARED`, or re-create your password file before re-creating the control file.

**See Also:** *Oracle8i Reference* for more information about the `REMOTE_LOGIN_PASSWORDFILE` parameter.

### Keywords and Parameters

---

<code>REUSE</code>	specifies that existing control files identified by the initialization parameter <code>CONTROL_FILES</code> can be reused, thus ignoring and overwriting any information they may currently contain. If you omit this clause and any of these control files already exists, Oracle returns an error.
<code>DATABASE</code> <i>database</i>	specifies the name of the database. The value of this parameter must be the existing database name established by the previous <code>CREATE DATABASE</code> statement or <code>CREATE CONTROLFILE</code> statement.
<code>SET DATABASE</code> <i>database</i>	changes the name of the database. The name of a database can be as long as eight bytes.

---

LOGFILE	<p>specifies the redo log files for your database. You must list all members of all redo log file groups.</p> <p><b>See Also:</b> "filespec" on page 7-516 for the syntax of <i>filespec</i>.</p> <p>GROUP <i>integer</i> specifies logfile group. If you specify GROUP values, Oracle verifies these values with the GROUP values when the database was last open.</p>
RESETLOGS	<p>ignores the contents of the files listed in the LOGFILE clause. These files do not have to exist. Each <i>filespec</i> in the LOGFILE clause must specify the SIZE parameter. Oracle assigns all online redo log file groups to thread 1 and enables this thread for public use by any instance. After using this clause, you must open the database using the RESETLOGS clause of the ALTER DATABASE statement.</p>
NORESETLOGS	<p>specifies that all files in the LOGFILE clause should be used as they were when the database was last open. These files must exist and must be the current online redo log files rather than restored backups. Oracle reassigns the redo log file groups to the threads to which they were previously assigned and reenables the threads as they were previously enabled.</p>
DATAFILE	<p>specifies the datafiles of the database. You must list all datafiles. These files must all exist, although they may be restored backups that require media recovery. See the syntax description of <i>filespec</i> in "filespec" on page 7-516.</p>
MAXLOGFILES <i>integer</i>	<p>specifies the maximum number of online redo log file groups that can ever be created for the database. Oracle uses this value to determine how much space in the control file to allocate for the names of redo log files. The default and maximum values depend on your operating system. The value that you specify should not be less than the greatest GROUP value for any redo log file group.</p>
MAXLOGMEMBERS <i>integer</i>	<p>specifies the maximum number of members, or identical copies, for a redo log file group. Oracle uses this value to determine how much space in the control file to allocate for the names of redo log files. The minimum value is 1. The maximum and default values depend on your operating system.</p>
MAXLOGHISTORY <i>integer</i>	<p>specifies the maximum number of archived redo log file groups for automatic media recovery of the Oracle Parallel Server. Oracle uses this value to determine how much space in the control file to allocate for the names of archived redo log files. The minimum value is 0. The default value is a multiple of the MAXINSTANCES value and depends on your operating system. The maximum value is limited only by the maximum size of the control file. This parameter is useful only if you are using Oracle with the Parallel Server option in both parallel mode and archive log mode.</p>
MAXDATAFILES <i>integer</i>	<p>specifies the initial sizing of the datafiles section of the control file at CREATE DATABASE or CREATE CONTROLFILE time. An attempt to add a file whose number is greater than MAXDATAFILES, but less than or equal to DB_FILES, causes the Oracle control file to expand automatically so that the datafiles section can accommodate more files.</p> <p>The number of datafiles accessible to your instance is also limited by the initialization parameter DB_FILES.</p>

## CREATE CONTROLFILE

---

---

<b>MAXINSTANCES</b> <i>integer</i>	specifies the maximum number of instances that can simultaneously have the database mounted and open. This value takes precedence over the value of the initialization parameter <code>INSTANCES</code> . The minimum value is 1. The maximum and default values depend on your operating system.
<b>ARCHIVELOG</b>	establishes the mode of archiving the contents of redo log files before reusing them. This clause prepares for the possibility of media recovery as well as instance or crash recovery.
<b>NOARCHIVELOG</b>	If you omit both the <code>ARCHIVELOG</code> clause and <code>NOARCHIVELOG</code> clause, Oracle chooses <code>NOARCHIVELOG</code> mode by default. After creating the control file, you can change between <code>ARCHIVELOG</code> mode and <code>NOARCHIVELOG</code> mode with the <code>ALTER DATABASE</code> statement.
<b>CHARACTER SET</b> <i>character_set</i>	<p>optionally reconstructs character set information in the control file. In case media recovery of the database is required, this information will be available before the database is open, so that tablespace names can be correctly interpreted during recovery. This clause is useful only if you are using a character set other than the default <code>US7ASCII</code>.</p> <p>If you are re-creating your control file and you are using Recovery Manager for tablespace recovery, and if you specify a different character set from the one stored in the data dictionary, then tablespace recovery will not succeed. (However, at database open, the control file character set will be updated with the correct character set from the data dictionary.)</p> <p><b>See Also:</b> <i>Oracle8i Recovery Manager User's Guide and Reference</i> for more information on tablespace recovery.</p> <hr/> <p><b>Note:</b> You cannot modify the character set of the database with this clause.</p> <hr/>

## Example

This statement re-creates a control file. In this statement, database `ORDERS_2` was created with the `F7DEC` character set.

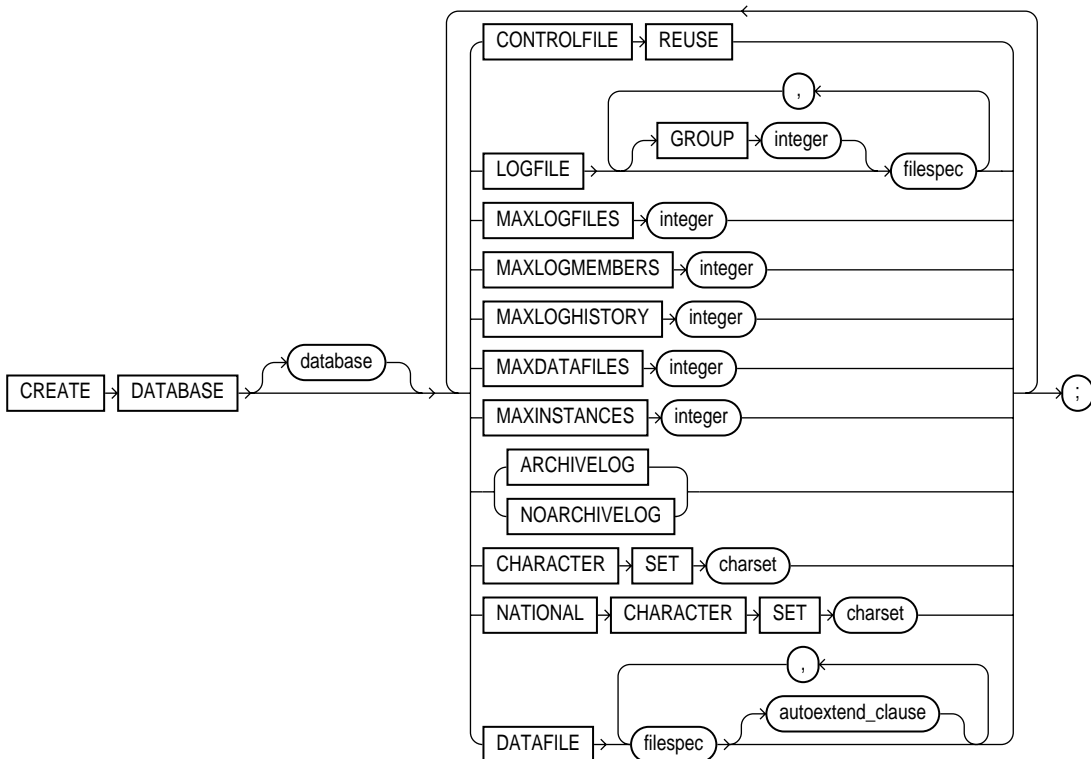
```
CREATE CONTROLFILE REUSE
  DATABASE orders_2
  LOGFILE GROUP 1 ('diskb:log1.log', 'diskc:log1.log') SIZE 50K,
            GROUP 2 ('diskb:log2.log', 'diskc:log2.log') SIZE 50K
  NORESETLOGS
  DATAFILE 'diska:dbone.dat' SIZE 2M
                                     MAXLOGFILES 5
                                     MAXLOGHISTORY 100
                                     MAXDATAFILES 10
                                     MAXINSTANCES 2
                                     ARCHIVELOG



  CHARACTER SET F7DEC;
```

# CREATE DATABASE

## Syntax

**WARNING:** This statement prepares a database for initial use and erases any data currently in the specified files. Use this statement only when you understand its ramifications.



**autoextend\_clause::=****maxsize\_clause::=**

**filespec:** See ["filespec"](#) on page 7-516.

**Purpose**

To create a database, making it available for general use.

This statement erases all data in any specified datafiles that already exist in order to prepare them for initial database use. If you use the statement on an existing database, all data in the datafiles is lost.

After creating the database, this statement mounts it in either exclusive or parallel mode (depending on the value of the `PARALLEL_SERVER` initialization parameter) and opens it, making it available for normal use. You can then create tablespaces and rollback segments for the database. For information on these tasks, see ["CREATE ROLLBACK SEGMENT"](#) on page 4-367 and ["CREATE TABLESPACE"](#) on page 4-419.

**See Also:**

- ["ALTER DATABASE"](#) on page 7-6 for information on modifying a database.
- *Oracle8i Java Developer's Guide* for information on creating an Oracle8i Java virtual machine.



## Prerequisites

You must have the OSDBA role enabled.

If the `REMOTE_LOGIN_PASSWORDFILE` initialization parameter is set to `EXCLUSIVE`, Oracle returns an error when you attempt to re-create the database. To avoid this message, either set the parameter to `SHARED`, or re-create your password file before re-creating the database.

**See Also:** *Oracle8i Reference* for more information about the `REMOTE_LOGIN_PASSWORDFILE` parameter.

## Keyword and Parameters

---

<code>database</code>	<p>is the name of the database to be created and can be up to 8 bytes long. The database name can contain only ASCII characters. Oracle writes this name into the control file. If you subsequently issue an <code>ALTER DATABASE</code> statement that explicitly specifies a database name, Oracle verifies that name with the name in the control file. Database names should also adhere to the rules described in "<a href="#">Schema Object Naming Rules</a>" on page 2-72.</p> <hr/> <p><b>Note:</b> You cannot use special characters from European or Asian character sets in a database name. For example, characters with umlauts are not allowed.</p> <hr/> <p>If you omit the database name from a <code>CREATE DATABASE</code> statement, Oracle uses the name specified by the initialization parameter <code>DB_NAME</code>. If the <code>DB_NAME</code> initialization parameter has been set, and you specify a different name from the value of that parameter, Oracle returns an error.</p>
<code>CONTROLFILE REUSE</code>	<p>reuses existing control files identified by the initialization parameter <code>CONTROL_FILES</code>, thus ignoring and overwriting any information they currently contain. Normally you use this clause only when you are re-creating a database, rather than creating one for the first time. You cannot use this clause if you also specify a parameter value that requires that the control file be larger than the existing files. These parameters are <code>MAXLOGFILES</code>, <code>MAXLOGMEMBERS</code>, <code>MAXLOGHISTORY</code>, <code>MAXDATAFILES</code>, and <code>MAXINSTANCES</code>.</p> <p>If you omit this clause and any of the files specified by <code>CONTROL_FILES</code> already exist, Oracle returns an error.</p>
<code>LOGFILE</code>	<p>specifies one or more files to be used as redo log files. Each <i>filespec</i> specifies a redo log file group containing one or more redo log file members (copies). For the syntax of <i>filespec</i>, see "<a href="#">filespec</a>" on page 7-516. All redo log files specified in a <code>CREATE DATABASE</code> statement are added to redo log thread number 1.</p>

---

	<p><code>GROUP <i>integer</i></code> uniquely identifies a redo log file group and can range from 1 to the value of the <code>MAXLOGFILES</code> parameter. A database must have at least two redo log file groups. You cannot specify multiple redo log file groups having the same <code>GROUP</code> value. If you omit this parameter, Oracle generates its value automatically. You can examine the <code>GROUP</code> value for a redo log file group through the dynamic performance table <code>V\$LOG</code>.</p> <p>If you omit the <code>LOGFILE</code> clause, Oracle creates two redo log file groups by default. The names and sizes of the default files depend on your operating system.</p>
<code>MAXLOGFILES <i>integer</i></code>	<p>specifies the maximum number of redo log file groups that can ever be created for the database. Oracle uses this value to determine how much space in the control file to allocate for the names of redo log files. The default, minimum, and maximum values depend on your operating system.</p>
<code>MAXLOGMEMBERS <i>integer</i></code>	<p>specifies the maximum number of members, or copies, for a redo log file group. Oracle uses this value to determine how much space in the control file to allocate for the names of redo log files. The minimum value is 1. The maximum and default values depend on your operating system.</p>
<code>MAXLOGHISTORY <i>integer</i></code>	<p>specifies the maximum number of archived redo log files for automatic media recovery with Oracle Parallel Server. Oracle uses this value to determine how much space in the control file to allocate for the names of archived redo log files. The minimum value is 0. The default value is a multiple of the <code>MAXINSTANCES</code> value and depends on your operating system. The maximum value is limited only by the maximum size of the control file.</p> <hr/> <p><b>Note:</b> This parameter is useful only if you are using Oracle with the Parallel Server option in parallel mode, and archivelog mode enabled.</p>
<code>MAXDATAFILES <i>integer</i></code>	<p>specifies the initial sizing of the datafiles section of the control file at <code>CREATE DATABASE</code> or <code>CREATE CONTROLFILE</code> time. An attempt to add a file whose number is greater than <code>MAXDATAFILES</code>, but less than or equal to <code>DB_FILES</code>, causes the Oracle control file to expand automatically so that the datafiles section can accommodate more files.</p> <p>The number of datafiles accessible to your instance is also limited by the initialization parameter <code>DB_FILES</code>.</p>
<code>MAXINSTANCES <i>integer</i></code>	<p>specifies the maximum number of instances that can simultaneously have this database mounted and open. This value takes precedence over the value of initialization parameter <code>INSTANCES</code>. The minimum value is 1. The maximum and default values depend on your operating system.</p>
<code>ARCHIVELOG</code>	<p>specifies that the contents of a redo log file group must be archived before the group can be reused. This clause prepares for the possibility of media recovery.</p>
<code>NOARCHIVELOG</code>	<p>specifies that the contents of a redo log file group need not be archived before the group can be reused. This clause does not allow for the possibility of media recovery.</p>

---

	The default is NOARCHIVELOG mode. After creating the database, you can change between ARCHIVELOG mode and NOARCHIVELOG mode with the ALTER DATABASE statement.										
CHARACTER SET	<p>specifies the character set the database uses to store data. The supported character sets and default value of this parameter depend on your operating system.</p> <p><b>Restriction:</b> You cannot specify any fixed-width multibyte character sets as the database character set.</p> <p><b>See Also:</b> <i>Oracle8i National Language Support Guide</i> for more information about character sets.</p>										
NATIONAL CHARACTER SET	specifies the national character set used to store data in columns specifically defined as NCHAR, NCLOB, or NVARCHAR2. If not specified, the national character set defaults to the database character set. See <i>Oracle8i National Language Support Guide</i> for valid character set names.										
DATAFILE	<p>specifies one or more files to be used as datafiles. See the syntax description of <i>filespec</i> in "filespec" on page 7-516. All these files become part of the SYSTEM tablespace. If you omit this clause, Oracle creates one datafile by default. The name and size of this default file depend on your operating system.</p> <hr/> <p><b>Note:</b> Oracle recommends that the total initial space allocated for the SYSTEM tablespace be a minimum of 5 megabytes.</p>										
<i>autoextend_clause</i>	<p>enables or disables the automatic extension of a datafile. If you do not specify this clause, datafiles are not automatically extended.</p> <hr/> <table> <tr> <td>OFF</td> <td>disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in ALTER DATABASE AUTOEXTEND or ALTER TABLESPACE AUTOEXTEND statements.</td> </tr> <tr> <td>ON</td> <td>enables autoextend.</td> </tr> <tr> <td>NEXT</td> <td>specifies the size in bytes of the next increment of disk space to be allocated to the datafile automatically when more extents are required. Use K or M to specify this size in kilobytes or megabytes. The default is the size of one data block.</td> </tr> <tr> <td>MAXSIZE</td> <td>specifies the maximum disk space allowed for automatic extension of the datafile.</td> </tr> <tr> <td>UNLIMITED</td> <td>sets no limit on the allocation of disk space to the datafile.</td> </tr> </table>	OFF	disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in ALTER DATABASE AUTOEXTEND or ALTER TABLESPACE AUTOEXTEND statements.	ON	enables autoextend.	NEXT	specifies the size in bytes of the next increment of disk space to be allocated to the datafile automatically when more extents are required. Use K or M to specify this size in kilobytes or megabytes. The default is the size of one data block.	MAXSIZE	specifies the maximum disk space allowed for automatic extension of the datafile.	UNLIMITED	sets no limit on the allocation of disk space to the datafile.
OFF	disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in ALTER DATABASE AUTOEXTEND or ALTER TABLESPACE AUTOEXTEND statements.										
ON	enables autoextend.										
NEXT	specifies the size in bytes of the next increment of disk space to be allocated to the datafile automatically when more extents are required. Use K or M to specify this size in kilobytes or megabytes. The default is the size of one data block.										
MAXSIZE	specifies the maximum disk space allowed for automatic extension of the datafile.										
UNLIMITED	sets no limit on the allocation of disk space to the datafile.										

---

## Examples

The following statement creates a small database using defaults for all arguments:

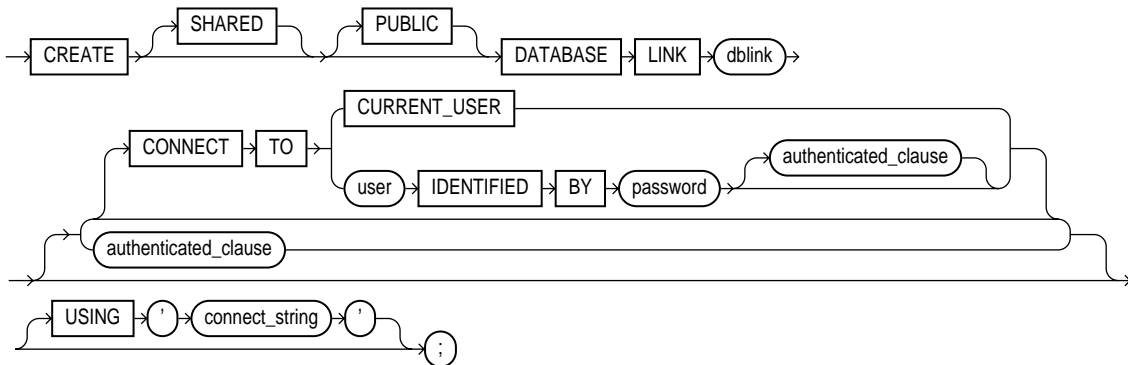
```
CREATE DATABASE ;
```

The following statement creates a database and fully specifies each argument:


```
CREATE DATABASE newtest
  CONTROLFILE REUSE
  LOGFILE
    GROUP 1 ('diskb:log1.log', 'diskc:log1.log') SIZE 50K,
    GROUP 2 ('diskb:log2.log', 'diskc:log2.log') SIZE 50K
  MAXLOGFILES 5
  MAXLOGHISTORY 100
  DATAFILE 'diska:dbone.dat' SIZE 2M
  MAXDATAFILES 10
  MAXINSTANCES 2
  ARCHIVELOG
  CHARACTER SET US7ASCII
  NATIONAL CHARACTER SET JA16SJISFIXED
  DATAFILE
    'disk1:df1.dbf' AUTOEXTEND ON
    'disk2:df2.dbf' AUTOEXTEND ON NEXT 10M MAXSIZE UNLIMITED;
```

## CREATE DATABASE LINK

### Syntax



#### **authenticated\_clause::=**



### Purpose

To create a database link. A **database link** is a schema object in the local database that allows you to access objects on a remote database. The remote database need not be an Oracle system.

Once you have created a database link, you can use it to refer to tables and views on the remote database. You can refer to a remote table or view in a SQL statement by appending *@dblink* to the table or view name. You can query a remote table or view with the SELECT statement. If you are using Oracle with the distributed option, you can also access remote tables and views using any of the following statements:

- "DELETE" on page 7-464
- "INSERT" on page 7-539
- "LOCK TABLE" on page 7-547
- "UPDATE" on page 7-615

For information about accessing remote tables or views with PL/SQL functions, procedures, packages, and datatypes, see *Oracle8i Application Developer's Guide - Fundamentals*. For information on distributed database systems, see *Oracle8i Distributed Database Systems*.

## Prerequisites

To create a private database link, you must have `CREATE DATABASE LINK` system privilege. To create a public database link, you must have `CREATE PUBLIC DATABASE LINK` system privilege.

You must have `CREATE SESSION` privilege on the remote Oracle database.

Net8 must be installed on both the local and remote Oracle databases.

To access non-Oracle systems you must use the Oracle Heterogeneous Services.

## Keyword and Parameters

---

<code>SHARED</code>	uses a single network connection to create a public database link that can be shared between multiple users. This clause is available only with the multi-threaded server configuration.  <b>See Also:</b> <i>Oracle8i Distributed Database Systems</i> for more information about shared database links.
<code>PUBLIC</code>	creates a public database link available to all users. If you omit this clause, the database link is private and is available only to you.
<i>dblink</i>	is the complete or partial name of the database link. For guidelines for naming database links, see " <a href="#">Referring to Objects in Remote Databases</a> " on page 2-79.  <b>Restrictions:</b> <ul style="list-style-type: none"><li>■ You cannot create a database link in another user's schema, and you cannot qualify <i>dblink</i> with the name of a schema. (Periods are permitted in names of database links, so Oracle interprets the entire name, such as <code>RALPH.LINKTOSALES</code>, as the name of a database link in your schema rather than as a database link named <code>LINKTOSALES</code> in the schema <code>RALPH</code>.)</li><li>■ The number of different database links that can appear in a single statement is limited to the value of the initialization parameter <code>OPEN_LINKS</code>.</li></ul>
<code>CONNECT TO</code>	enables a connection to the remote database.
<code>CURRENT_USER</code>	creates a <b>current user database link</b> . The current user must be a global user with a valid account on the remote database for the link to succeed.  If the database link is used directly, that is, not from within a stored object, then the current user is the same as the connected user.

When executing a stored object (such as a procedure, view, or trigger) that initiates a database link, `CURRENT_USER` is the username that owns the stored object, and not the username that called the object. For example, if the database link appears inside procedure `SCOTT.P` (created by `SCOTT`), and user `JANE` calls procedure `SCOTT.P`, the current user is `SCOTT`.

However, if the stored object is an invoker-rights function, procedure, or package, the invoker's authorization ID is used to connect as a remote user. For example, if the privileged database link appears inside procedure `SCOTT.P` (an invoker-rights procedure created by `SCOTT`), and user `JANE` calls procedure `SCOTT.P`, then `CURRENT_USER` is `JANE` and the procedure executes with `JANE`'s privileges.

**See Also:** "[CREATE FUNCTION](#)" on page 7-284 for more information on invoker-rights functions.

<code>user IDENTIFIED BY password</code>	is the username and password used to connect to the remote database (fixed user database link). If you omit this clause, the database link uses the username and password of each user who is connected to the database (connected user database link).
<code>authenticated_clause</code>	specifies the username and password on the target instance. This clause authenticates the user to the remote server and is required for security. The specified username and password must be a valid username and password on the remote instance. The username and password are used only for authentication. No other operations are performed on behalf of this user.
<code>USING 'connect string'</code>	You must specify this clause when using the <code>SHARED</code> clause. specifies the service name of a remote database. For information on specifying remote databases, see <i>Net8 Administrator's Guide</i> .

## Examples

**CURRENT\_USER Example** The following statement defines a current-user database link:

```
CREATE DATABASE LINK sales.hq.acme.com
  CONNECT TO CURRENT_USER
  USING 'sales';
```

**Fixed User Example** The following statement defines a fixed-user database link named `SALES.HQ.ACME.COM`:

```
CREATE DATABASE LINK sales.hq.acme.com
  CONNECT TO scott IDENTIFIED BY tiger
  USING 'sales';
```

Once this database link is created, you can query tables in the schema `SCOTT` on the remote database in this manner:

```
SELECT *  
FROM emp@sales.hq.acme.com;
```

You can also use DML statements to modify data on the remote database:

```
INSERT INTO accounts@sales.hq.acme.com(acc_no, acc_name, balance)  
VALUES (5001, 'BOWER', 2000);
```

```
UPDATE accounts@sales.hq.acme.com  
SET balance = balance + 500;
```

```
DELETE FROM accounts@sales.hq.acme.com  
WHERE acc_name = 'BOWER';
```

You can also access tables owned by other users on the same database. This statement assumes SCOTT has access to ADAM's DEPT table:

```
SELECT *  
FROM adams.dept@sales.hq.acme.com;
```

The previous statement connects to the user SCOTT on the remote database and then queries ADAM's DEPT table.

A synonym may be created to hide the fact that SCOTT's EMP table is on a remote database. The following statement causes all future references to EMP to access a remote EMP table owned by SCOTT:

```
CREATE SYNONYM emp  
FOR scott.emp@sales.hq.acme.com;
```


**PUBLIC Example** The following statement defines a shared public fixed user database link named SALES.HQ.ACME.COM that refers to user SCOTT with password TIGER on the database specified by the string service name 'SALES':

```
CREATE SHARED PUBLIC DATABASE LINK sales.hq.acme.com  
CONNECT TO scott IDENTIFIED BY tiger  
AUTHENTICATED BY anupam IDENTIFIED BY bhide  
USING 'sales';
```




# CREATE DIMENSION


## Syntax




### level\_clause::=




### hierarchy\_clause::=



### join\_clause::=



**attribute\_clause::=**



## Purpose

To create a **dimension**. A dimension defines a parent-child relationship between pairs of column sets, where all the columns of a column set must come from the same table. However, columns in one column set (or "level") can come from a different table than columns in another set. The optimizer uses these relationships with materialized views to perform **query rewrite**. The Summary Advisor uses these relationships to recommend creation of specific materialized views.

---

**Note:** Oracle does not automatically validate the relationships you declare when creating a dimension. To validate the relationships specified in the *hierarchy\_clause* and the *join\_clause*, you must run the `DBMS_OLAP.VALIDATE_DIMENSION` procedure. For information on this procedure, see *Oracle8i Supplied PL/SQL Packages Reference*.

---

### See Also:

- ["CREATE MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-318 for more information on materialized views.
- *Oracle8i Data Warehousing Guide* for more information on query rewrite, the optimizer and the Summary Advisor.

## Prerequisites

To create a dimension in your own schema, you must have the `CREATE DIMENSION` system privilege. To create a dimension in another user's schema, you must have the `CREATE ANY DIMENSION` system privilege. In either case, you must have the `SELECT` object privilege on any objects referenced in the dimension.

## Keywords and Parameters

---

<i>schema</i>	is the schema in which the dimension will be created. If you do not specify <i>schema</i> , Oracle creates the dimension in your own schema.								
<i>dimension</i>	is the name of the dimension. The name must be unique within its schema.								
<i>level_clause</i>	defines a level in the dimension. A level defines dimension hierarchies and attributes. <table> <tr> <td><i>level</i></td> <td>is the name of the level</td> </tr> <tr> <td><i>level_table . level_column</i></td> <td>specifies from the columns in the level. You can specify up to 32 columns. The tables you specify in this clause must already exist.</td> </tr> <tr> <td colspan="2"><b>Restrictions:</b></td> </tr> <tr> <td colspan="2"> <ul style="list-style-type: none"> <li>■ All of the columns in a level must come from the same table.</li> <li>■ If columns in different levels come from different tables, then you must specify the <i>join_clause</i>.</li> <li>■ The set of columns you specify must be unique to this level.</li> <li>■ The columns you specify cannot be specified in any other dimension.</li> <li>■ Each <i>level_column</i> must be non-null. (However, these columns need not have NOT NULL constraints.)</li> </ul> </td> </tr> </table>	<i>level</i>	is the name of the level	<i>level_table . level_column</i>	specifies from the columns in the level. You can specify up to 32 columns. The tables you specify in this clause must already exist.	<b>Restrictions:</b>		<ul style="list-style-type: none"> <li>■ All of the columns in a level must come from the same table.</li> <li>■ If columns in different levels come from different tables, then you must specify the <i>join_clause</i>.</li> <li>■ The set of columns you specify must be unique to this level.</li> <li>■ The columns you specify cannot be specified in any other dimension.</li> <li>■ Each <i>level_column</i> must be non-null. (However, these columns need not have NOT NULL constraints.)</li> </ul>	
<i>level</i>	is the name of the level								
<i>level_table . level_column</i>	specifies from the columns in the level. You can specify up to 32 columns. The tables you specify in this clause must already exist.								
<b>Restrictions:</b>									
<ul style="list-style-type: none"> <li>■ All of the columns in a level must come from the same table.</li> <li>■ If columns in different levels come from different tables, then you must specify the <i>join_clause</i>.</li> <li>■ The set of columns you specify must be unique to this level.</li> <li>■ The columns you specify cannot be specified in any other dimension.</li> <li>■ Each <i>level_column</i> must be non-null. (However, these columns need not have NOT NULL constraints.)</li> </ul>									
<i>hierarchy_clause</i>	defines a linear hierarchy of levels in the dimension. Each hierarchy forms a chain of parent-child relationships among the levels in the dimension. Hierarchies in a dimension are independent of each other. They may (but need not) have columns in common. <p>Each level in the dimension should be specified at most once in this clause, and each level must already have been named in the <i>level_clause</i>.</p> <table> <tr> <td><i>hierarchy</i></td> <td>is the name of the hierarchy. This name must be unique in the dimension.</td> </tr> <tr> <td><i>child_level</i></td> <td>is the name of a level that has an <i>n</i>:1 relationship with a parent level: the <i>level_columns</i> of <i>child_level</i> cannot be null, and each <i>child_level</i> value uniquely determines the value of the next named <i>parent_level</i>.           <p>If the child <i>level_table</i> is different from the parent <i>level_table</i>, you must specify a join relationship between them in the <i>join_clause</i>.</p> </td> </tr> <tr> <td><i>parent_level</i></td> <td>is the name of a level.</td> </tr> </table>	<i>hierarchy</i>	is the name of the hierarchy. This name must be unique in the dimension.	<i>child_level</i>	is the name of a level that has an <i>n</i> :1 relationship with a parent level: the <i>level_columns</i> of <i>child_level</i> cannot be null, and each <i>child_level</i> value uniquely determines the value of the next named <i>parent_level</i> . <p>If the child <i>level_table</i> is different from the parent <i>level_table</i>, you must specify a join relationship between them in the <i>join_clause</i>.</p>	<i>parent_level</i>	is the name of a level.		
<i>hierarchy</i>	is the name of the hierarchy. This name must be unique in the dimension.								
<i>child_level</i>	is the name of a level that has an <i>n</i> :1 relationship with a parent level: the <i>level_columns</i> of <i>child_level</i> cannot be null, and each <i>child_level</i> value uniquely determines the value of the next named <i>parent_level</i> . <p>If the child <i>level_table</i> is different from the parent <i>level_table</i>, you must specify a join relationship between them in the <i>join_clause</i>.</p>								
<i>parent_level</i>	is the name of a level.								

<i>join_clause</i>	<p>specifies an inner equijoin relationship for a dimension whose columns are contained in multiple tables. This clause is required and permitted only when the columns specified in the hierarchy are not all in the same table.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ The <i>child_key_columns</i> must be non-null and the parent key must be unique and non-null. You need not define constraints to enforce these conditions, but queries may return incorrect results if these conditions are not true.</li><li>■ Each child key must join with a key in the <i>parent_level</i> table.</li><li>■ Self-joins are not permitted (that is, the <i>child_key_columns</i> cannot be in the same table as <i>parent_level</i>).</li></ul>
<i>child_key_column</i>	<p>specifies one or more columns that are join-compatible with columns in the parent level.</p> <p>If you do not specify the schema and table of each <i>child_column</i>, the schema and table are inferred from the <code>CHILD OF</code> relationship in the <i>hierarchy_clause</i>. If you do specify the schema and column of a <i>child_key_column</i>, the schema and table must match the schema and table of columns that comprise the child of <i>parent_level</i> in the <i>hierarchy_clause</i>.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ All of the child-key columns must come from the same table.</li><li>■ The number of child-key columns must match the number of columns in <i>parent_level</i>, and the columns must be joinable.</li><li>■ Do not specify multiple child key columns unless the parent level consists of multiple columns.</li></ul> <p>You can specify only one <i>join_clause</i> for a given pair of levels in the same hierarchy.</p>
<i>parent_level</i>	<p>is the name of a level.</p>
<i>attribute_clause</i>	<p>specifies the columns that are uniquely determined by a hierarchy level. The columns in <i>level</i> must all come from the same table as the <i>dependent_columns</i>. The <i>dependent_columns</i> need not have been specified in the <i>level_clause</i>.</p> <p>For example, if the hierarchy levels are <i>city</i>, <i>state</i>, and <i>country</i>, then <i>city</i> might determine <i>mayor</i>, <i>state</i> might determine <i>governor</i>, and <i>country</i> might determine <i>president</i>.</p>

---

## Examples

This statement creates a `TIME` dimension on table `TIME_TAB`, and creates a `GEOG` dimension on tables `CITY`, `STATE`, and `COUNTRY`.

```
CREATE DIMENSION time
  LEVEL curDate          IS time_tab.curDate
  LEVEL month            IS time_tab.month
```


```
LEVEL qtr                IS time_tab.qtr
LEVEL year               IS time_tab.year
LEVEL fiscal_week       IS time_tab.fiscal_week
LEVEL fiscal_qtr        IS time_tab.fiscal_qtr
LEVEL fiscal_year       IS time_tab.fiscal_year
HIERARCHY month_rollup (
    curDate              CHILD OF
    month                CHILD OF
    qtr                  CHILD OF
    year)
HIERARCHY fiscal_year_rollup (
    curDate              CHILD OF
    fiscal_week          CHILD OF
    fiscal_qtr           CHILD OF
    fiscal_year )
ATTRIBUTE curDate       DETERMINES (holiday, dayOfWeek)
ATTRIBUTE month         DETERMINES (yr_ago_month, qtr_ago_month)
ATTRIBUTE fiscal_qtr    DETERMINES yr_ago_qtr
ATTRIBUTE year          DETERMINES yr_ago ;

CREATE DIMENSION geog
LEVEL cityID             IS (city.city, city.state)
LEVEL stateID           IS state.state
LEVEL countryID         IS country.country
HIERARCHY political_rollup (
    cityID              CHILD OF
    stateID             CHILD OF
    countryID)
JOIN KEY city.state REFERENCES stateID
JOIN KEY state.country REFERENCES countryID);
```

---

## CREATE DIRECTORY

### Syntax



### Purpose

To create a directory object. A directory object specifies an alias for a directory on the server's file system where external binary file LOBs (`BFILEs`) are located. You can use directory names when referring to `BFILEs` in your PL/SQL code and OCI calls, rather than hard-coding the operating system pathname, thereby allowing greater file management flexibility.

All directories are created in a single namespace and are not owned by an individual's schema. You can secure access to the `BFILEs` stored within the directory structure by granting object privileges on the directories to specific users.

When you create a directory, you are automatically granted the `READ` object privilege and can grant `READ` privileges to other users and roles. The DBA can also grant this privilege to other users and roles.

#### See Also:

- ["Large Object \(LOB\) Datatypes"](#) on page 2-23 for more information on `BFILE` objects.
- ["GRANT object\\_privileges"](#) on page 7-532 for more information on granting object privileges.

### Prerequisites

You must have `CREATE ANY DIRECTORY` system privileges to create directories.

You must also create a corresponding operating system directory for file storage. Your system or database administrator must ensure that the operating system directory has the correct read permissions for Oracle processes.

Privileges granted for the directory are created independently of the permissions defined for the operating system directory. Therefore, the two may or may not correspond exactly. For example, an error occurs if user `SCOTT` is granted `READ`

privilege on the directory schema object, but the corresponding operating system directory does not have `READ` permission defined for Oracle processes.

## Keywords and Parameters

<code>OR REPLACE</code>	<p>re-creates the directory database object if it already exists. You can use this clause to change the definition of an existing directory without dropping, re-creating, and regrating database object privileges previously granted on the directory.</p> <p>Users who had previously been granted privileges on a redefined directory can still access the directory without being regranted the privileges</p> <p>For information on removing a directory from the database, see "<a href="#">DROP DIRECTORY</a>" on page 7-477.</p>
<i>directory</i>	<p>is the name of the directory object to be created. The maximum length of <i>directory</i> is 30 bytes. You cannot qualify a directory object with a schema name.</p> <p><b>Note:</b> Oracle does not verify that the directory you specify actually exists. Therefore, take care that you specify a valid directory in your operating system. In addition, if your operating system uses case-sensitive pathnames, be sure you specify the directory in the correct format. (However, you need not include a trailing slash at the end of the pathname.)</p>
<i>'path_name'</i>	<p>is the full pathname of the operating system directory on the server where the files are located. The single quotes are required, with the result that the path name is case sensitive.</p>

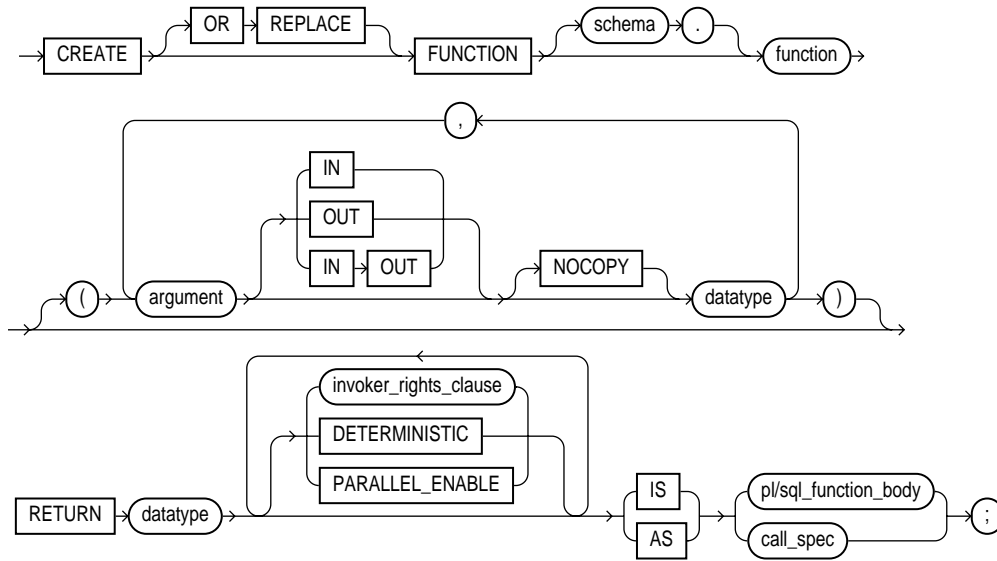
## Example

The following statement redefines directory database object `BFILE_DIR` to enable access to `BFILEs` stored in the operating system directory `/PRIVATE1/LOB/FILES:`


```
CREATE OR REPLACE DIRECTORY bfile_dir AS '/privatel/LOB/files';
```

# CREATE FUNCTION


## Syntax




### invoker\_rights\_clause::=




### call\_spec::=



### Java\_declaration::=





**C\_declaration::=****Purpose**

To create a stored function or a call specification.

A **stored function** (also called a **user function**) is a set of PL/SQL statements you can call by name. Stored functions are very similar to procedures, except that a function returns a value to the environment in which it is called. User functions can be used as part of a SQL expression. For a general discussion of procedures and functions, see ["CREATE PROCEDURE"](#) on page 7-353. For examples of creating functions, see ["Examples"](#) on page 7-289.

A **call specification** declares a Java method or a third-generation language (3GL) routine so that it can be called from SQL and PL/SQL. The call specification tells Oracle which Java method, or which named function in which shared library, to invoke when a call is made. It also tells Oracle what type conversions to make for the arguments and return value.

The `CREATE FUNCTION` statement creates a function as a standalone schema object. You can also create a function as part of a package. For information on creating packages, see ["CREATE PACKAGE"](#) on page 7-344.

For information on modifying a function, see ["ALTER FUNCTION"](#) on page 7-28. For information on shared libraries, see ["CREATE LIBRARY"](#) on page 7-316. For information on dropping a standalone function, see ["DROP FUNCTION"](#) on page 7-478.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals* for more information about registering external functions.

**Prerequisites**

Before a stored function can be created, the user SYS must run the SQL script `DBMSSTD.X.SQL`. The exact name and location of this script depend on your operating system.

To create a function in your own schema, you must have the `CREATE PROCEDURE` system privilege. To create a function in another user's schema, you must have the

CREATE ANY PROCEDURE system privilege. To replace a function in another user's schema, you must have the ALTER ANY PROCEDURE system privilege.

To invoke a call specification, you may need additional privileges (for example, EXECUTE privileges on C library for a C call specification).

To embed a CREATE FUNCTION statement inside an Oracle precompiler program, you must terminate the statement with the keyword END-EXEC followed by the embedded SQL statement terminator for the specific language.

**See Also:** *PL/SQL User's Guide and Reference* or *Oracle8i Java Stored Procedures Developer's Guide* for more information on such prerequisites.

## Keywords and Parameters

---

<code>OR REPLACE</code>	re-creates the function if it already exists. Use this clause to change the definition of an existing function without dropping, re-creating, and regrating object privileges previously granted on the function. If you redefine a function, Oracle recompiles it. For information on recompiling functions, see " <a href="#">ALTER FUNCTION</a> " on page 7-28.  Users who had previously been granted privileges on a redefined function can still access the function without being regranted the privileges.  If any function-based indexes depend on the function, Oracle marks the indexes <code>DISABLED</code> .
<i>schema</i>	is the schema to contain the function. If you omit <i>schema</i> , Oracle creates the function in your current schema.
<i>function</i>	is the name of the function to be created. If creating the function results in compilation errors, Oracle returns an error. You can see the associated compiler error messages with the <code>SHOW ERRORS</code> command.

### Restrictions on User-Defined Functions

User-defined functions cannot be used in situations that require an unchanging definition. Thus, you cannot use user-defined functions:

- In a CHECK constraint clause of a CREATE TABLE or ALTER TABLE statement
- In a DEFAULT clause of a CREATE TABLE or ALTER TABLE statement

---

In addition, when a function is called from within a query or DML statement, the function cannot:

- Have `OUT` or `IN OUT` parameters
- Commit or roll back the current transaction, create or roll back to a savepoint, or alter the session or the system. DDL statements implicitly commit the current transaction, so a user-defined function cannot execute any DDL statements.
- Write to the database, if the function is being called from a `SELECT` statement. However, a function called from a subquery in a DML statement can write to the database.
- Write to the same table that is being modified by the statement from which the function is called, if the function is called from a DML statement.

Except for the restriction on `OUT` and `IN OUT` parameters, Oracle enforces these restrictions not only for the function called directly from the SQL statement, but also for any functions that function calls, and on any functions called from the SQL statements executed by that function or any function it calls.

<i>argument</i>	is the name of an argument to the function. If the function does not accept arguments, you can omit the parentheses following the function name.
<code>IN</code>	specifies that you must supply a value for the argument when calling the function. This is the default.
<code>OUT</code>	specifies the function will set the value of the argument.
<code>IN OUT</code>	specifies that a value for the argument can be supplied by you and may be set by the function.
<code>NOCOPY</code>	<p>instructs Oracle to pass this argument as fast as possible. This clause can significantly enhance performance when passing a large value like a record, an index-by table, or a varray to an <code>OUT</code> or <code>IN OUT</code> parameter. (<code>IN</code> parameter values are always passed <code>NOCOPY</code>.)</p> <ul style="list-style-type: none"> <li>■ When you specify <code>NOCOPY</code>, assignments made to a package variable may show immediately in this parameter (or assignments made to this parameter may show immediately in a package variable) if the package variable is passed as the actual assignment corresponding to this parameter.</li> <li>■ Similarly, changes made either to this parameter or to another parameter may be visible immediately through both names if the same variable is passed to both.</li> <li>■ If the function is exited with an unhandled exception, any assignment made to this parameter may be visible in the caller's variable.</li> </ul> <p>These effects may or may not occur on any particular call. You should use <code>NOCOPY</code> only when these effects would not matter.</p>
<i>datatype</i>	<p>is the datatype of an argument. An argument can have any datatype supported by PL/SQL.</p> <p>The datatype cannot specify a length, precision, or scale. Oracle derives the length, precision, or scale of an argument from the environment from which the function is called.</p>

<code>RETURN datatype</code>	<p>specifies the datatype of the function's return value. Because every function must return a value, this clause is required. The return value can have any datatype supported by PL/SQL.</p> <p>The datatype cannot specify a length, precision, or scale. Oracle derives the length, precision, or scale of the return value from the environment from which the function is called. For information on PL/SQL datatypes, see <i>PL/SQL User's Guide and Reference</i>.</p>
<code>invoker_rights_clause</code>	<p>lets you specify whether the function executes with the privileges and in the schema of the user who owns it or with the privileges and in the schema of <code>CURRENT_USER</code>. For information on how <code>CURRENT_USER</code> is determined, see <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i>.</p> <p>This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the function.</p> <p><b>See Also:</b> <i>PL/SQL User's Guide and Reference</i>.</p>
<code>AUTHID CURRENT_USER</code>	<p>specifies that the function executes with the privileges of <code>CURRENT_USER</code>. This clause creates an "invoker-rights function."</p> <p>This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of <code>CURRENT_USER</code>. External names in all other statements resolve in the schema in which the function resides.</p>
<code>AUTHID DEFINER</code>	<p>specifies that the function executes with the privileges of the owner of the schema in which the function resides, and that external names resolve in the schema where the function resides. This is the default.</p>
<code>DETERMINISTIC</code>	<p>is an optimization hint that allows the system to use a saved copy of the function's return result (if such a copy is available). The saved copy could come from a materialized view, a function-based index, or a redundant call to the same function in the same SQL statement. The query optimizer can choose whether to use the saved copy or re-call the function.</p> <p>The function should reliably return the same result value whenever it is called with the same values for its arguments. Therefore, do not define the function to use package variables or to access the database in any way that might affect the function's return result, because the results of doing so will not be captured if the system chooses not to call the function.</p> <p>A function must be declared <code>DETERMINISTIC</code> in order to be called in the expression of a function-based index, or from the query of a materialized view if that view is marked <code>REFRESH FAST</code> or <code>ENABLE QUERY REWRITE</code>.</p> <p>For information on materialized views, see <i>Oracle8i Data Warehousing Guide</i>. For information on function-based indexes, see "<a href="#">CREATE INDEX</a>" on page 7-291.</p>
<code>PARALLEL_ENABLE</code>	<p>is an optimization hint indicating that the function can be executed from a parallel execution server of a parallel query operation. The function should not use session state, such as package variables, as those variables may not be shared among the parallel execution servers.</p> <p><b>See Also:</b> <i>Oracle8i Application Developer's Guide - Fundamentals</i>.</p>

---

<i>pl/sql_</i> <i>subprogram_body</i>	declares the function in a PL/SQL subprogram body. <b>See Also:</b> <i>Oracle8i Application Developer's Guide - Fundamentals</i> for more information on PL/SQL subprograms.
<i>call_spec</i>	maps a Java or C method name, parameter types, and return type to their SQL counterparts. In <i>Java_declaration</i> , 'string' identifies the Java implementation of the method. <b>See Also:</b> <ul style="list-style-type: none"> <li>■ <i>Oracle8i Java Stored Procedures Developer's Guide</i>.</li> <li>■ <i>Oracle8i Application Developer's Guide - Fundamentals</i> for an explanation of the parameters and semantics of the <i>C_declaration</i>.</li> </ul>
AS EXTERNAL	is an alternative way of declaring a C method. This clause has been deprecated and is supported for backward compatibility only. Oracle Corporation recommends that you use the AS LANGUAGE C syntax.

---

## Examples

The following statement creates the function GET\_BAL.

```
CREATE FUNCTION get_bal(acc_no IN NUMBER)
RETURN NUMBER
IS acc_bal NUMBER(11,2);
BEGIN
    SELECT balance
    INTO acc_bal
    FROM accounts
    WHERE account_id = acc_no;
RETURN(acc_bal);
END;
```

The GET\_BAL function returns the balance of a specified account.

When you call the function, you must specify the argument ACC\_NO, the number of the account whose balance is sought. The datatype of ACC\_NO is NUMBER.

The function returns the account balance. The RETURN clause of the CREATE FUNCTION statement specifies the datatype of the return value to be NUMBER.

The function uses a SELECT statement to select the BALANCE column from the row identified by the argument ACC\_NO in the ACCOUNTS table. The function uses a RETURN statement to return this value to the environment in which the function is called.

The function created above can be used in a SQL statement. For example:


```
SELECT get_bal(100) FROM DUAL;
```

The following statement creates PL/SQL standalone function `GET_VAL` that registers the C routine `C_GET_VAL` as an external function. (The parameters have been omitted from this example.)


```
CREATE FUNCTION get_val
  ( x_val IN NUMBER,
    y_val IN NUMBER,
    image IN LONG RAW )
RETURN BINARY_INTEGER AS LANGUAGE C
  NAME "c_get_val"
  LIBRARY c_utils
  PARAMETERS (...);
```

# CREATE INDEX

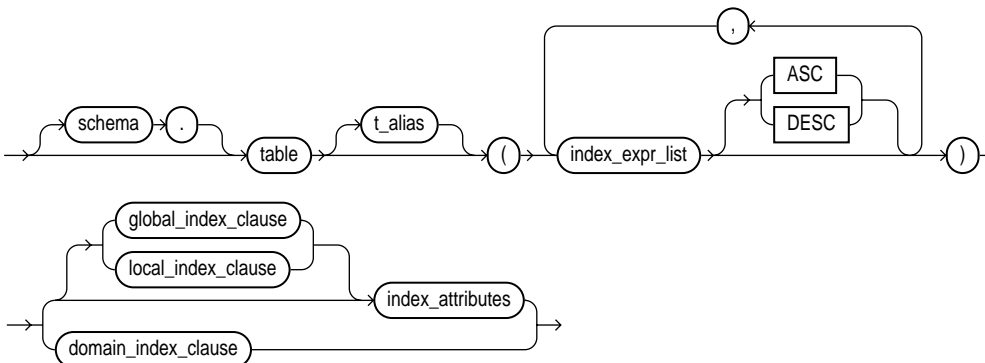
## Syntax




### cluster\_index\_clause::=



### table\_index\_clause::=

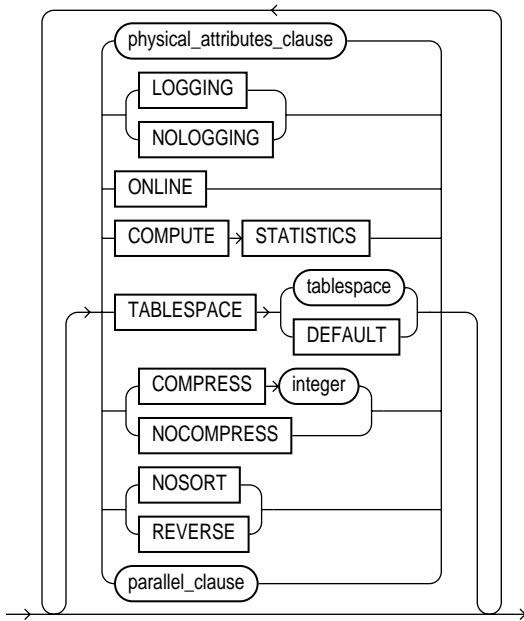


### index\_expr\_list::=




# CREATE INDEX


## index\_attributes::=




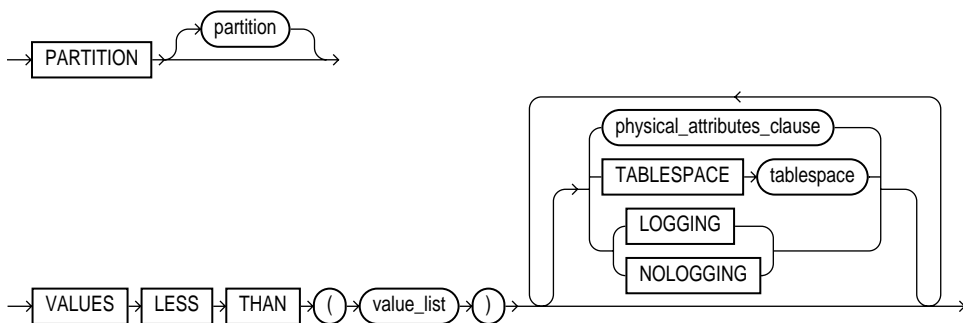


## physical\_attributes\_clause::=



## domain\_index\_clause::=






**global\_index\_clause::=****global\_partition\_clause::=****local\_index\_clauses::=****on\_range\_partitioned\_table\_clause::=**

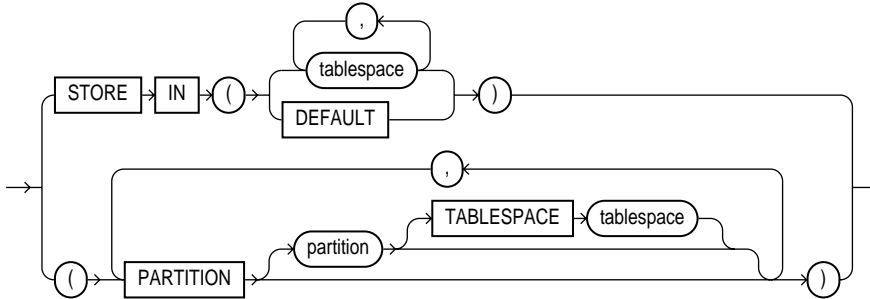
## CREATE INDEX

---

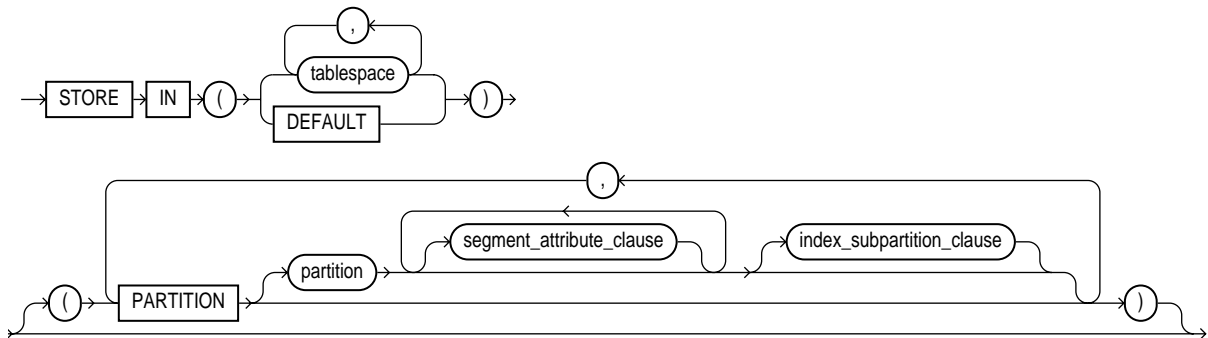
### segment\_attributes\_clause::=

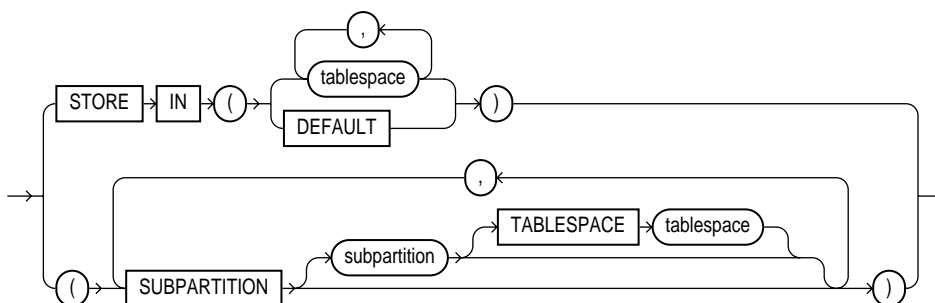



### on\_hash\_partitioned\_table\_clause::=



### on\_composite\_partitioned\_table\_clause::=



**index\_subpartition\_clause::=****parallel\_clause::=**

**storage\_clause:** See ["storage\\_clause"](#) on page 7-605.

**Purpose**

To create an index on

- One or more columns of a table, a partitioned table, an index-organized table, or a cluster
- One or more scalar typed object attributes of a table or a cluster
- A nested table storage table for indexing a nested table column

To create a **domain index**, which is an instance of an application-specific index of type *indextype*.

An **index** is a schema object that contains an entry for each value that appears in the indexed column(s) of the table or cluster and provides direct, fast access to rows. A **partitioned index** consists of partitions containing an entry for each value that appears in the indexed column(s) of the table. A **function-based index** is an index on expressions. It enables you to construct queries that evaluate the value returned by an expression, which in turn may include functions (built-in or user-defined).

For a discussion of indexes, see *Oracle8i Concepts*. For information on modifying an index, see "[ALTER INDEX](#)" on page 7-30.

## Prerequisites

To create an index in your own schema, one of the following conditions must be true:

- The table or cluster to be indexed must be in your own schema.
- You must have `INDEX` privilege on the table to be indexed.
- You must have `CREATE ANY INDEX` system privilege.

To create an index in another schema, you must have `CREATE ANY INDEX` system privilege. Also, the owner of the schema to contain the index must have either space quota on the tablespaces to contain the index or index partitions, or `UNLIMITED TABLESPACE` system privilege.

To create a domain index in your own schema, you must also have `EXECUTE` privilege on the indextype. If you are creating a domain index in another user's schema, the index owner also must have `EXECUTE` privilege on the indextype and its underlying implementation type. Before creating a domain index, you should first define the indextype. See "[CREATE INDEXTYPE](#)" on page 7-309.

To create a function-based index in your own schema on your own table, you must have the `QUERY REWRITE` system privilege. To create the index in another schema or on another schema's table, you must have the `GLOBAL QUERY REWRITE` privilege. The table owner must also have the `EXECUTE` object privilege on the function(s) used in the function-based index.

## Keywords and Parameters

---

### UNIQUE

specifies that the value of the column (or columns) upon which the index is based must be unique. If the index is local nonprefixed (see *local\_index\_clause* below), then the index key must contain the partitioning key.

Oracle recommends that you do not explicitly define `UNIQUE` indexes on tables. Uniqueness is strictly a logical concept and should be associated with the *definition* of a table. Therefore, define `UNIQUE` integrity constraints on the desired columns.

**See Also:** The "[constraint\\_clause](#)" on page 7-233.

#### Restrictions:

- You cannot specify both `UNIQUE` and `BITMAP`.
- You cannot specify `UNIQUE` for a domain index.

---

<code>BITMAP</code>	<p>specifies that <i>index</i> is to be created as a bitmap, rather than as a B-tree. Bitmap indexes store the rowids associated with a key value as a bitmap. Each bit in the bitmap corresponds to a possible rowid, and if the bit is set, it means that the row with the corresponding rowid contains the key value. The internal representation of bitmaps is best suited for applications with low levels of concurrent transactions, such as data warehousing. See <i>Oracle8i Concepts</i> and <i>Oracle8i Designing and Tuning for Performance</i> for more information about using bitmap indexes.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot specify <code>BITMAP</code> when creating a global partitioned index or an index-organized table.</li><li>■ You cannot specify both <code>UNIQUE</code> and <code>BITMAP</code>.</li><li>■ You cannot specify <code>BITMAP</code> for a domain index.</li></ul>
<code>schema</code>	<p>is the schema to contain the index. If you omit <i>schema</i>, Oracle creates the index in your own schema.</p>
<code>index</code>	<p>is the name of the index to be created. An <i>index</i> can contain several partitions.</p>
<code>cluster_index_clause</code>	<p>specifies the cluster for which a cluster index is to be created. If you do not qualify <i>cluster</i> with <i>schema</i>, Oracle assumes the cluster is in your current schema. You cannot create a cluster index for a hash cluster.</p> <p><b>See Also:</b> "<a href="#">CREATE CLUSTER</a>" on page 7-254.</p>
<code>table_index_clause</code>	<p>specifies <i>table</i> (and its attributes) on which you are defining the index. If you do not qualify <i>table</i> with <i>schema</i>, Oracle assumes the table is contained in your own schema.</p> <p>You create an index on a nested table column by creating the index on the nested table storage table. Include the <code>NESTED_TABLE_ID</code> pseudocolumn of the storage table to create a <code>UNIQUE</code> index, which effectively ensures that the rows of a nested table value are distinct.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ If the index is local, then <i>table</i> must be partitioned.</li><li>■ If the table is index-organized, this statement creates a secondary index. You cannot specify <code>BITMAP</code> or <code>REVERSE</code> for this secondary index, and the combined size of the index key and the logical rowid should be less than half the block size.</li><li>■ If <i>table</i> is a temporary table, the index will also be temporary with the same scope (session or transaction) as <i>table</i>. The following restrictions apply to indexes on temporary table:<ul style="list-style-type: none"><li>- The index cannot be a partitioned index or a domain index.</li><li>- You cannot specify the <i>physical_attributes_clause</i> or the <i>parallel_clause</i>.</li><li>- You cannot specify <code>LOGGING</code>, <code>NOLOGGING</code>, or <code>TABLESPACE</code>.</li></ul></li></ul> <p><b>See Also:</b> "<a href="#">CREATE TABLE</a>" on page 4-381 and <i>Oracle8i Concepts</i> for more information on temporary tables.</p>

<i>t_alias</i>	specifies a correlation name (alias) for the table upon which you are building the index. <b>Note:</b> This alias is required if the <i>index_expression_list</i> references any object type attributes or object type methods. See " <a href="#">Function-based Index on Type Method Example</a> " on page 7-306.
<i>index_expr_list</i>	lets you specify the column or column expression upon which the index is based.
<i>column</i>	is the name of a column in the table. A bitmap index can have a maximum of 30 columns. Other indexes can have as many as 32 columns. <b>Restriction:</b> You cannot create an index on columns or attributes whose type is user-defined, LONG, LONG RAW, LOB, or REF, except that Oracle supports an index on REF type columns or attributes that have been defined with a SCOPE clause.  You can create an index on a scalar object attribute column or on the system-defined NESTED_TABLE_ID column of the nested table storage table. If you specify an object attribute column, the column name must be qualified with the table name. If you specify a nested table column attribute, it must be qualified with the outermost table name, the containing column name, and all intermediate attribute names leading to the nested table column attribute.
<i>column_expression</i>	is an expression built from columns of <i>table</i> , constants, SQL functions, and user-defined functions. When you specify <i>column_expression</i> , you create a <b>function-based index</b> .  Name resolution of the function is based on the schema of the index creator. User-defined functions used in <i>column_expression</i> are fully name resolved during the CREATE INDEX operation.  After creating a function-based index, collect statistics on both the index and its base table using the ANALYZE statement (see " <a href="#">ANALYZE</a> " on page 7-200). Oracle cannot use the function-based index until these statistics have been generated.

---

**Notes on function-based indexes:**

- When you subsequently query a table that uses a function-based index, you must ensure in the query that *column\_expression* is not null. However, Oracle will use a function-based index in a query even if the columns specified in the WHERE clause are in a different order than their order in the *column\_expression* that defined the function-based index.  
**See Also:** The [Function-Based Index Example](#) on page 7-305.
- If the function on which the index is based becomes invalid or is dropped, Oracle marks the index DISABLED. Queries on a DISABLED index fail if the optimizer chooses to use the index. DML operations on a DISABLED index fail unless the index is also marked UNUSABLE **and** the parameter SKIP\_UNUSABLE\_INDEXES is set to true.  
**See Also:** "[ALTER SESSION](#)" on page 7-83 for more information on this parameter).
- Oracle's use of function-based indexes is also affected by the setting of the QUERY\_REWRITE\_ENABLED session parameter.  
**See Also:** "[ALTER SESSION](#)" on page 7-83.

- If a public synonym for a function, package, or type is used in *column\_expression*, and later an actual object with the same name is created in the table owner's schema, then Oracle will disable the function-based index. When you subsequently enable the function-based index using `ALTER INDEX ... ENABLE` or `ALTER INDEX ... REBUILD`, the function, package, or type used in the *column\_expression* will continue to resolve to the function, package, or type to which the public synonym originally pointed. It will not resolve to the new function, package, or type.
- If the definition of a function-based index generates internal conversion to character data, use caution when changing NLS parameter settings. Function-based indexes use the current database settings for NLS parameters. If you reset these parameters at the session level, queries using the function-based index may return incorrect results. Two exceptions are the collation parameters (`NLS_SORT` and `NLS_COMP`). Oracle handles the conversions correctly even if these have been reset at the session level.

---

#### Restrictions on function-based indexes:

- Any user-defined function referenced in *column\_expression* must be `DETERMINISTIC`.
- For a function-based globally partitioned index, the *column\_expression* cannot be the partitioning key.
- All functions must be specified with parentheses, even if they have no parameters. Otherwise Oracle interprets them as column names.
- Any function you specify in *column\_expression* must return a repeatable value. For example, you cannot specify the `SYSDATE` or `USER` function or the `ROWNUM` pseudocolumn.
- You cannot build a function-based index on `LOB`, `REF`, nested table, or varray columns. In addition, the function in *column\_expression* cannot take as arguments any objects with attributes of type `LOB`, `REF`, nested table, or varray.
- The *column\_expression* cannot contain any aggregate functions.

**See Also:** "CREATE FUNCTION" on page 7-284 and *PL/SQL User's Guide and Reference*.

ASC | DESC

specifies whether the index should be created in ascending or descending order. Indexes on character data are created in ascending or descending order of the character values in the database character set.

Oracle treats descending indexes as if they were function-based indexes. You do not need the `QUERY REWRITE` or `GLOBAL QUERY REWRITE` privileges to create them, as you do with other function-based indexes. However, as with other function-based indexes, Oracle does not use descending indexes until you first analyze the index and the table on which the index is defined. See the *column\_expression* clause of this statement.

**Restriction:** You cannot specify either of these clauses for a domain index. You cannot specify `DESC` for a reverse index. Oracle ignores `DESC` if index is bitmapped or if the `COMPATIBLE` initialization parameter is set to a value less than 8.1.

**index\_attributes**

*physical\_attributes\_clause* establishes values for physical and storage characteristics for the index. See "CREATE TABLE" on page 4-381.

**Restriction:** You cannot specify the `PCTUSED` parameter for an index.

`PCTFREE` is the percentage of space to leave free for updates and insertions within each of the index's data blocks.

*storage\_clause* establishes the storage characteristics for the index. See the "storage\_clause" on page 7-605.

`TABLESPACE` is the name of the tablespace to hold the index, index partition, or index subpartition. If you omit this clause, Oracle creates the index in the default tablespace of the owner of the schema containing the index.

For a local index, you can specify the keyword `DEFAULT` in place of *tablespace*. New partitions or subpartitions added to the local index will be created in the same tablespace(s) as the corresponding partitions or subpartitions of the underlying table.

`COMPRESS` enables key compression, which eliminates repeated occurrence of key column values and may substantially reduce storage. Use *integer* to specify the prefix length (number of prefix columns to compress).

- For unique indexes, the valid range of prefix length values is from 1 to the number of key columns minus 1. The default prefix length is the number of key columns minus 1.
- For nonunique indexes, the valid range of prefix length values is from 1 to the number of key columns. The default prefix length is the number of key columns.

Oracle compresses only nonpartitioned indexes that are nonunique or unique indexes of at least two columns.

**Restriction:** You cannot specify `COMPRESS` for a bitmapped index.

`NOCOMPRESS` disables key compression. This is the default.

`NOSORT` indicates to Oracle that the rows are stored in the database in ascending order, so that Oracle does not have to sort the rows when creating the index. If the rows of the indexed column or columns are not stored in ascending order, Oracle returns an error. For greatest savings of sort time and space, use this clause immediately after the initial load of rows into a table.

**Restrictions:**

- You cannot specify `REVERSE` with this clause.
- You cannot use this clause to create a cluster, partitioned, or bitmap index.
- You cannot specify this clause for a secondary index on an index-organized table.

`REVERSE` stores the bytes of the index block in reverse order, excluding the rowid. You cannot specify `NOSORT` with this clause.

You cannot reverse a bitmap index or an index-organized table.



---

LOGGING |  
NOLOGGING

specifies that the creation of the index will be logged (LOGGING) or not logged (NOLOGGING) in the redo log file. It also specifies that subsequent Direct Loader (SQL\*Loader) and direct-load INSERT operations against the index are logged or not logged. LOGGING is the default.

If *index* is nonpartitioned, this is the logging attribute of the index.

If *index* is partitioned, the logging attribute specified is

- The default value of all partitions specified in the CREATE statement (unless you specify LOGGING | NOLOGGING in the PARTITION description clause)
- The default value for the segments associated with the index partitions
- The default value for local index partitions or subpartitions added implicitly during subsequent ALTER TABLE ... ADD PARTITION operations

In NOLOGGING mode, data is modified with minimal logging (to mark new extents INVALID and to record dictionary changes). When applied during media recovery, the extent invalidation records mark a range of blocks as logically corrupt, since the redo data is not logged. Thus if you cannot afford to lose this index, it is important to take a backup after the NOLOGGING operation.

If the database is run in ARCHIVELOG mode, media recovery from a backup taken before the LOGGING operation will re-create the index. However, media recovery from a backup taken before the NOLOGGING operation will not re-create the index.

The logging attribute of the index is independent of that of its base table.

If you omit this clause, the logging attribute is that of the tablespace in which it resides.

**See Also:** *Oracle8i Concepts* and *Oracle8i Parallel Server Concepts* for more information about logging and parallel DML.

ONLINE

specifies that DML operations on the table will be allowed during creation of the index. For a description of online index building and rebuilding, see *Oracle8i Concepts*.

**Restriction:** Parallel DML is not supported during online index building. If you specify ONLINE and then issue parallel DML statements, Oracle returns an error.

COMPUTE  
STATISTICS

enables you to collect statistics at relatively little cost during the creation of an index. These statistics are stored in the data dictionary for ongoing use by the optimizer in choosing a plan of execution for SQL statements.

The types of statistics collected depend on the type of index you are creating.

---

**Note:** If you create an index using another index (instead of a table), the original index might not provide adequate statistical information. Therefore, Oracle generally uses the base table to compute the statistics, which will improve the statistics but may negatively affect performance.

---

Additional methods of collecting statistics are available in PL/SQL packages and procedures.

**See Also:** *Oracle8i Supplied PL/SQL Packages Reference*.

---

<i>global_index_clause</i>	<p>specifies that the partitioning of the index is user defined and is not equipartitioned with the underlying table. By default, nonpartitioned indexes are global indexes.</p>
PARTITION BY RANGE	<p>specifies that the global index is partitioned on the ranges of values from the columns specified in <i>column_list</i>. You cannot specify this clause for a local index.</p>
( <i>column_list</i> )	<p>is the name of the column(s) of a table on which the index is partitioned. The <i>column_list</i> must specify a left prefix of the index column list.</p> <p>You cannot specify more than 32 columns in <i>column_list</i>, and the columns cannot contain the ROWID pseudocolumn or a column of type ROWID.</p> <hr/> <p><b>Note:</b> If your enterprise has or will have databases using different character sets, use caution when partitioning on character columns. The sort sequence of characters is not identical in all character sets.</p> <p><b>See Also:</b> <i>Oracle8i National Language Support Guide</i> for more information on character set support.</p> <hr/>
PARTITION <i>partition</i>	<p>describes the individual partitions. The number of clauses determines the number of partitions. If you omit <i>partition</i>, Oracle generates a name with the form SYS_P<i>n</i>.</p>
VALUES LESS THAN ( <i>value_list</i> )	<p>specifies the (noninclusive) upper bound for the current partition in a global index. The <i>value_list</i> is a comma-separated, ordered list of literal values corresponding to <i>column_list</i> in the <i>partition_by_range_clause</i>. Always specify MAXVALUE as the <i>value_list</i> of the last partition.</p> <hr/> <p><b>Note:</b> If <i>index</i> is partitioned on a DATE column, and if the NLS date format does not specify the century with the year, you must use the TO_DATE function with a 4-character format mask for the year. The NLS date format is determined implicitly by NLS_TERRITORY or explicitly by NLS_DATE_FORMAT.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ <i>Oracle8i National Language Support Guide</i> for more information on these initialization parameters.</li> <li>■ The "<a href="#">Partitioned Table Example</a>" on page 4-414.</li> </ul> <hr/> <p><b>Restriction:</b> You cannot specify this clause for a local index.</p> <hr/>
<i>local_index_clauses</i>	<p>specify that the index is partitioned on the same columns, with the same number of partitions and the same partition bounds as <i>table</i>. Oracle automatically maintains LOCAL index partitioning as the underlying table is repartitioned.</p>
<i>on_range_partitioned_table_clause</i>	<p>describes an index on a range-partitioned table.</p>

---

PARTITION <i>partition</i>	<p>describes the individual partitions. The number of clauses determines the number of partitions. For a local index, the number of index partitions must be equal to the number of the table partitions, and in the same order.</p> <p>If you omit <i>partition</i>, Oracle generates a name that is consistent with the corresponding table partition. If the name conflicts with an existing index partition name, the form <code>SYS_Pn</code> is used.</p>
<i>on_hash_partitioned_table_clause</i>	<p>describes an index on a hash-partitioned table. If you do not specify <i>partition</i>, Oracle uses the name of the corresponding base table partition, unless it conflicts with an explicitly specified name of another index partition. In this case, Oracle generates a name of the form <code>SYS_Pnnn</code>.</p> <p>You can optionally specify <code>TABLESPACE</code> for all index partitions or for one or more individual partitions. If you do not specify <code>TABLESPACE</code> at the index or partition level, Oracle stores each index partition in the same tablespace as the corresponding table partition.</p>
<i>on_composite_partitioned_table_clause</i>	<p>describes an index on a composite-partitioned table. The first <code>STORE IN</code> clause specifies the default tablespace for the index subpartitions. You can override this storage by specifying a different tablespace in the <i>index_subpartitioning_clause</i>.</p> <p>If you do not specify <code>TABLESPACE</code> for subpartitions either in this clause or in the <i>index_subpartitioning_clause</i>, Oracle uses the tablespace specified for <i>index</i>. If you also do not specify <code>TABLESPACE</code> for <i>index</i>, Oracle stores the subpartition in the same tablespace as the corresponding table subpartition.</p>
STORE IN	lets you specify how index hash partitions (for a hash-partitioned index) or index subpartitions (for a composite-partitioned index) are to be distributed across various tablespaces. The number of tablespaces does not have to equal the number of index partitions. If the number of index partitions is greater than the number of tablespaces, Oracle cycles through the names of the tablespaces.
DEFAULT	is valid only for a local index on a hash or composite-partitioned table. This clause overrides any tablespace specified at the index level for a partition or subpartition, and stores the index partition or subpartition in the same partition as the corresponding table partition or subpartition.
<i>index_subpartition_clause</i>	specifies one or more tablespaces in which to store all subpartitions in <i>partition</i> or one or more individual subpartitions in <i>partition</i> . The subpartition inherits all other attributes from <i>partition</i> . Attributes not specified for <i>partition</i> are inherited from <i>index</i> .

---

<i>domain_index_clause</i>	<p>specifies that index is a domain index.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ The <i>index_expr_list</i> can specify only a single column.</li> <li>■ You can define only one domain index on a column.</li> <li>■ You cannot specify a BITMAP, UNIQUE, or function-based domain index.</li> <li>■ You cannot create a local domain index on a partitioned table.</li> <li>■ You cannot create a domain index on a partitioned table with row movement enabled.</li> </ul>
<i>column</i>	<p>specifies the table columns or object attributes on which the index is defined. Each <i>column</i> can have only one domain index defined on it.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot create a domain index on a column of datatype REF, varray, nested table, LONG, or LONG RAW.</li> <li>■ You can create a domain index on a column of user-defined type, but not on an attribute of a column of user-defined type if that attribute itself is a user-defined type.</li> </ul>
<i>indextype</i>	<p>specifies the name of the indextype. This name should be a valid schema object that you have already defined. See "<a href="#">CREATE INDEXTYPE</a>" on page 7-309.</p>
PARAMETERS 'string'	<p>specifies the parameter string that is passed uninterpreted to the appropriate indextype routine. The maximum length of the parameter string is 1000 characters.</p> <p>Once the domain index is created, Oracle invokes this routine (see <i>Oracle8i Data Cartridge Developer's Guide</i> for information on these routines.) If the routine does not return successfully, the domain index is marked FAILED. The only operation supported on an failed domain index is DROP INDEX.</p>
<i>parallel_clause</i>	<p>causes creation of the index to be parallelized. For additional information, see the <a href="#">Notes</a> to the <a href="#">parallel_clause</a> of "<a href="#">CREATE TABLE</a>" on page 4-381.</p>
NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the PARALLEL_THREADS_PER_CPU initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution servers. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .

---

## Examples

**PARALLEL Example** The following statement creates an index using 10 parallel execution servers, 5 to scan `SCOTT.EMP` and another 5 to populate the `EMP_IDX` index:

```
CREATE INDEX emp_idx
  ON scott.emp (ename)
  PARALLEL 5;
```

**COMPRESS Example** To create an index with the `COMPRESS` clause, you might issue the following statement:

```
CREATE INDEX emp_idx2 ON emp(job, ename) COMPRESS 1;
```

The index will compress repeated occurrences of `JOB` column values.

**NOLOGGING Example** To quickly create an index in parallel on a table that was created using a fast parallel load (so all rows are already sorted), you might issue the following statement. (Oracle will choose the appropriate degree of parallelism.)

```
CREATE INDEX i_loc
  ON big_table (akey)
  NOSORT
  NOLOGGING
  PARALLEL;
```

**Cluster Index Example** To create an index for the `EMPLOYEE` cluster, issue the following statement:

```
CREATE INDEX ic_emp ON CLUSTER employee;
```

No index columns are specified, because the index is automatically built on all the columns of the cluster key. For cluster indexes, all rows are indexed.

**NULL Example** Consider the following statement:

```
SELECT ename FROM emp WHERE comm IS NULL;
```

The above query does not use an index created on the `COMM` column unless it is a bitmap index.

**Function-Based Index Example** The following statements creates a function-based index on the `EMP` table based on an uppercase evaluation of the `ENAME` column:

```
CREATE INDEX emp_i ON emp (UPPER(ename));
```

To ensure that Oracle will use the index rather than performing a full table scan, be sure that the value of the function is not null in subsequent queries. For example, the statement

```
SELECT * FROM emp WHERE UPPER(ename) IS NOT NULL
      ORDER BY UPPER(ename);
```

is guaranteed to use the index, but without the `WHERE` clause, Oracle may perform a full table scan.

In the next statements showing index creation and subsequent query, Oracle will use index `EMP_FI` even though the columns are in reverse order in the query:

```
CREATE INDEX emp_fi ON emp(colb + cola);

SELECT * FROM emp WHERE colb + cola > 500;
```

**Function-based Index on Type Method Example** This example entails an object type `RECTANGLE` containing two number attributes: `length` and `width`. The `AREA()` method computes the area of the rectangle.

```
CREATE TYPE rectangle AS OBJECT
( length NUMBER,
  width NUMBER,
  MEMBER FUNCTION area RETURN NUMBER DETERMINISTIC
);

CREATE OR REPLACE TYPE BODY rectangle AS
  MEMBER FUNCTION area RETURN NUMBER IS
  BEGIN
    RETURN (length*width);
  END;
END;
```

Now, if you create a table `RECTTAB` of type `RECTANGLE`, you can create a function-based index on the `AREA()` method as follows:

```
CREATE TABLE recttab OF rectangle;
CREATE INDEX area_idx ON recttab x (x.area());
```

You can use this index efficiently to evaluate a query of the form:

```
SELECT * FROM recttab x WHERE x.area() > 100;
```

**Computing Statistics Example** The following statement collects statistics on the nonpartitioned EMP\_INDX index:

```
CREATE INDEX emp_idx ON emp(empno) COMPUTE STATISTICS;
```

The type of statistics collected depends on the type of index you are creating.

**See Also:** *Oracle8i Concepts*.

**Partitioned Index Example** The following statement creates a global prefixed index STOCK\_IX on table STOCK\_XACTIONS with two partitions, one for each half of the alphabet. The index partition names are system generated:

```
CREATE INDEX stock_idx ON stock_xactions
(stock_symbol, stock_series)
GLOBAL PARTITION BY RANGE (stock_symbol)
(PARTITION VALUES LESS THAN ('N') TABLESPACE ts3,
PARTITION VALUES LESS THAN (MAXVALUE) TABLESPACE ts4);
```

**Index on Hash-Partitioned Table Example.** This statement creates a local index on the ITEM column of the SALES table. The STORE IN clause immediately following LOCAL indicates that SALES is hash partitioned. Oracle will distribute the hash partitions between the TBS1 and TBS2 tablespaces:

```
CREATE INDEX sales_idx ON sales(item) LOCAL
STORE IN (tbs1, tbs2);
```

**Index on Composite-Partitioned Table Example.** This statement creates a local index on the SALES table, which is composite-partitioned. The STORAGE clause specifies default storage attributes for the index. The STORE IN clause specifies one or more default tablespaces for the index subpartitions. However, this default is overridden for the four subpartitions of partition Q3\_1997, because separate TABLESPACE is specified.

```
CREATE INDEX sales_idx ON sales(sale_date, item)
STORAGE (INITIAL 1M, MAXEXTENTS UNLIMITED)
LOCAL
STORE IN (tbs1, tbs2, tbs3, tbs4, tbs5)
(PARTITION q1_1997, PARTITION q2_1997,
PARTITION q3_1997
(SUBPARTITION q3_1997_s1 TABLESPACE ts2,
SUBPARTITION q3_1997_s2 TABLESPACE ts4,
SUBPARTITION q3_1997_s3 TABLESPACE ts6,
SUBPARTITION q3_1997_s4 TABLESPACE ts8),
PARTITION q4_1997,
```

```
PARTITION q1_1998);
```

**Bitmap Index Example** To create a bitmap partitioned index on a table with four partitions, issue the following statement:

```
CREATE BITMAP INDEX partno_ix
  ON lineitem(partno)
  TABLESPACE ts1
  LOCAL (PARTITION quarter1 TABLESPACE ts2,
         PARTITION quarter2 STORAGE (INITIAL 10K NEXT 2K),
         PARTITION quarter3 TABLESPACE ts2,
         PARTITION quarter4);
```

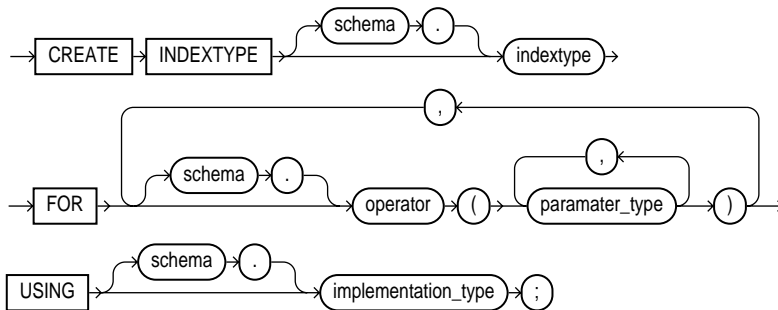
**Nested Table Example** In the following example, UNIQUE index UNIQ\_PROJ\_INDX is created on storage table NESTED\_PROJECT\_TABLE. Including pseudocolumn NESTED\_TABLE\_ID ensures distinct rows in nested table column PROJS\_MANAGED:

```
CREATE TYPE proj_type AS OBJECT
  (proj_num NUMBER, proj_name VARCHAR2(20));
CREATE TYPE proj_table_type AS TABLE OF proj_type;
CREATE TABLE employee ( emp_num NUMBER, emp_name CHAR(31),
  projs_managed proj_table_type )
  NESTED TABLE projs_managed STORE AS nested_project_table;
CREATE UNIQUE INDEX uniq_proj_indx
  ON nested_project_table ( NESTED_TABLE_ID, proj_num);
```



## CREATE INDEXTYPE

### Syntax



### Purpose

To create an **indextype**, which is an object that specifies the routines that manage a domain (application-specific) index. Indextypes reside in the same namespace as tables, views, and other schema objects. This statement binds the indextype name to an implementation type, which in turn specifies and refers to user-defined index functions and procedures that implement the indextype.

**See Also:** *Oracle8i Data Cartridge Developer's Guide* and *Oracle8i Concepts* for more information on implementing indextypes.

### Prerequisites

To create an indextype in your own schema, you must have the `CREATE INDEXTYPE` system privilege. To create an indextype in another schema, you must have `CREATE ANY INDEXTYPE` system privilege. In either case, you must have the `EXECUTE` object privilege on the implementation type and the supported operators.

An indextype supports one or more operators, so before creating an indextype, you should first design the operator or operators to be supported and provide functional implementation for those operators.

**See Also:** ["CREATE OPERATOR"](#) on page 7-339.

## Keywords and Parameters

---

<i>schema</i>	is the name of the schema in which the indextype resides. If you omit <i>schema</i> , Oracle creates the indextype in your own schema.
<i>indextype</i>	is the name of the indextype to be created.
FOR	specifies the list of operators supported by the indextype.
<i>schema</i>	is the schema containing the operator. If you omit <i>schema</i> , Oracle assumes the operator is in your own schema.
<i>operator</i>	specifies the name of the operator supported by the indextype.
<i>parameter_type</i>	lists the types of parameters to the operator.
	All the operators listed in this clause should be valid operators.
USING	specifies the type that provides the implementation for the new indextype.
<i>implementation_type</i>	is the name of the type that implements the appropriate Oracle Data Cartridge interface (ODCI). <ul style="list-style-type: none"><li>■ You must specify a valid type that implements the routines in the ODCI interface.</li><li>■ The implementation type must reside in the same schema as the indextype.</li></ul> For additional information on this interface, see <i>Oracle8i Data Cartridge Developer's Guide</i> .

---

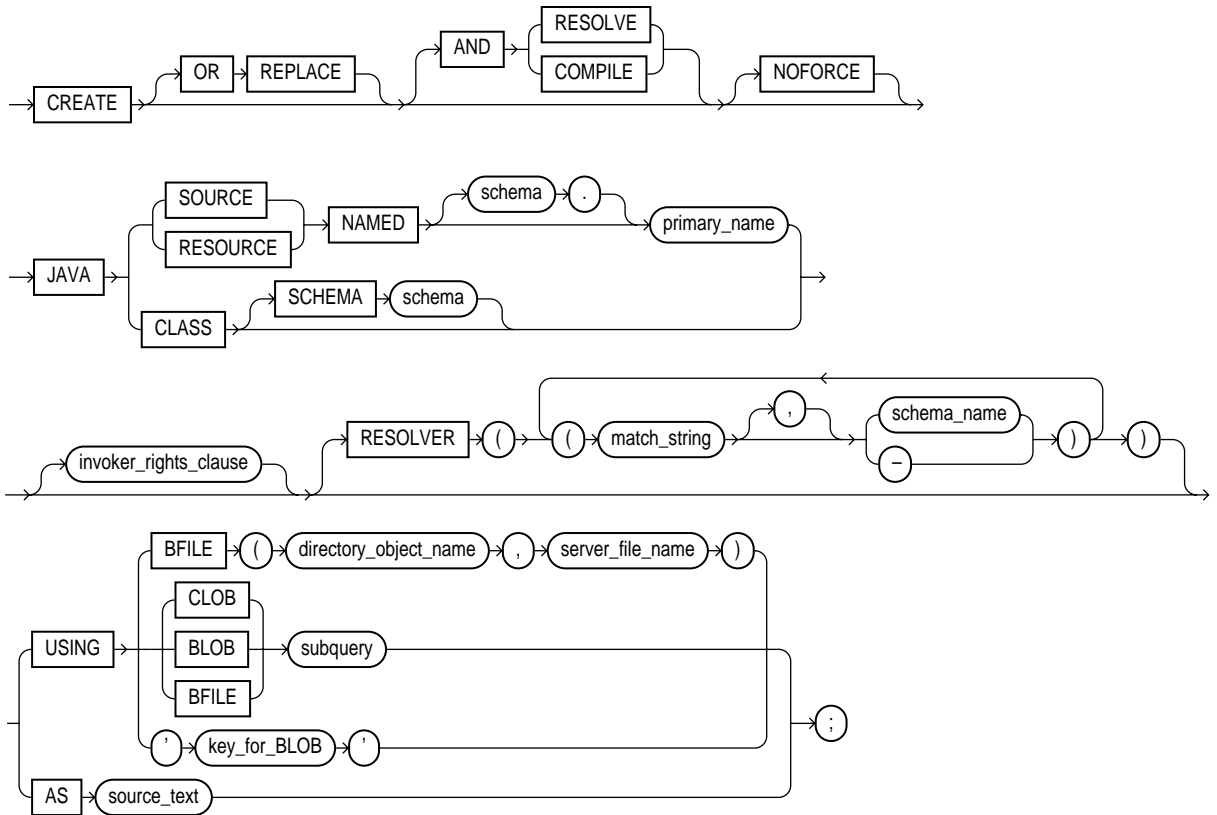
## Example

The following statement creates an indextype named `TextIndexType` and specifies the `CONTAINS` operator that is supported by the indextype and the `TextIndexMethods` type that implements the index interface:


```
CREATE INDEXTYPE TextIndexType
  FOR contains (VARCHAR2, VARCHAR2)
  USING TextIndexMethods;
```

# CREATE JAVA

## Syntax



**invoker\_rights\_clause::=**



## Purpose

To create a schema object containing a Java source, class, or resource. For information on the following topics, see these books:

- For Java concepts, see *Oracle8i Java Developer's Guide*.
- For Java stored procedures, see *Oracle8i Java Stored Procedures Developer's Guide*
- For SQLJ, see *Oracle8i SQLJ Developer's Guide and Reference*.
- For JDBC, see *Oracle8i JDBC Developer's Guide and Reference*.
- For CORBA and EJB, see *Oracle8i Enterprise JavaBeans and CORBA Developer's Guide*.

## Prerequisites

To create or replace a schema object containing a Java source, class, or resource in your own schema, you must have `CREATE PROCEDURE` system privilege. To create such a schema object in another user's schema, you must have `CREATE ANY PROCEDURE` system privilege. To replace such a schema object in another user's schema, you must also have `ALTER ANY PROCEDURE` system privilege.

## Keywords and Parameters

---

`OR REPLACE`

re-creates the schema object containing the Java class, source, or resource if it already exists. Use this clause to change the definition of an existing object without dropping, re-creating, and regranting object privileges previously granted.

If you redefine a Java schema object and specify `RESOLVE` or `COMPILE`, Oracle recompiles or resolves the object. Whether or not the resolution or compilation is successful, Oracle invalidates classes that reference the Java schema object. For additional information, see "[ALTER JAVA](#)" on page 7-45.

Users who had previously been granted privileges on a redefined function can still access the function without being regranted the privileges.

`RESOLVE` |  
`COMPILE`

are synonymous keywords. They specify that Oracle should attempt to resolve the Java schema object that is created if this statement succeeds.

- When applied to a class, resolution of referenced names to other class schema objects occurs.
- When applied to a source, source compilation occurs.

**Restriction:** You cannot specify this clause for a Java resource.

---

NOFORCE	rolls back the results of this CREATE command if you have specified either RESOLVE or COMPILE, and the resolution or compilation fails. If you do not specify this option, Oracle takes no action if the resolution or compilation fails (that is, the created schema object remains).
JAVA SOURCE	loads a Java source file.
JAVA CLASS	loads a Java class file.
JAVA RESOURCE	loads a Java resource file.
NAMED	<p>is <b>required</b> for a Java source or resource. The <i>primary_name</i> must be enclosed in double quotation marks.</p> <ul style="list-style-type: none"> <li>■ For a Java source, this clause specifies the name of the schema object in which the source code is held. A successful CREATE JAVA SOURCE statement will also create additional schema objects to hold each of the Java classes defined by the source.</li> <li>■ For a Java resource, this clause specifies the name of the schema object to hold the Java resource.</li> <li>■ Use double quotation marks to preserve lower- or mixed-case <i>primary_name</i>.</li> </ul> <p>If you do not specify <i>schema</i>, Oracle creates the object in your own schema.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify NAMED for a Java class.</li> <li>■ The <i>primary_name</i> cannot contain a database link.</li> </ul>
SCHEMA <i>schema</i>	applies only to a Java class. This optional clause specifies the schema in which the object containing the Java file resides. If you do not specify <i>schema</i> and you do not specify NAMED (above), Oracle creates the object in your own schema.
<i>invoker_rights_clause</i>	<p>specifies whether the methods of the class execute with the privileges and in the schema of the user who owns the class or with the privileges and in the schema of CURRENT_USER. For information on how CURRENT_USER is determined, see <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i>.</p> <p>This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the member functions and procedures of the type.</p> <p><b>See Also:</b> <i>Oracle8i Java Stored Procedures Developer's Guide</i>.</p>
AUTHID CURRENT_USER	<p>specifies that the methods of the class execute with the privileges of CURRENT_USER. This clause is the default and creates an "invoker-rights class."</p> <p>This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of CURRENT_USER. External names in all other statements resolve in the schema in which the methods reside.</p>

	AUTHID	specifies that the methods of the class execute with the privileges of the owner of the schema in which the class resides, and that external names resolve in the schema where the class resides.
	DEFINER	
RESOLVER		specifies a mapping of the fully qualified Java name to a Java schema object, where <ul style="list-style-type: none"><li>▪ <i>match_string</i> is either a fully qualified Java name, a wildcard that can match such a Java name, or a wildcard that can match any name.</li><li>▪ <i>schema_name</i> designates a schema to be searched for the corresponding Java schema object.</li><li>▪ A dash (-) as an alternative to <i>schema_name</i> indicates that if <i>match_string</i> matches a valid Java name, Oracle can leave the name unresolved. The resolution succeeds, but the name cannot be used at run time by the class.</li></ul> This mapping is stored with the definition of the schema objects created in this command for use in later resolutions (either implicit or in explicit ALTER ... RESOLVE statements).
USING		determines a sequence of character (CLOB or BFILE) or binary (BLOB or BFILE) data for the Java class or resource. Oracle uses the sequence of characters to define one file for a Java class or resource, or one source file and one or more derived classes for a Java source.
	BFILE	identifies a previously created file on the operating system ( <i>directory_object_name</i> ) and server file ( <i>server_file_name</i> ) containing the sequence. BFILE is usually interpreted as a character sequence by CREATE JAVA SOURCE and as a binary sequence by CREATE JAVA CLASS or CREATE JAVA RESOURCE.
	CLOB/BLOB/BFILE <i>subquery</i>	supplies a query that selects a single row and column of the type specified (CLOB, BLOB, or BFILE). The value of the column makes up the sequence of characters.
	<i>key_for_BLOB</i>	supplies the following implicit query: <pre>SELECT LOB FROM CREATE\$JAVA\$LOB\$TABLE WHERE NAME = 'key_for_BLOB';</pre> <b>Restriction:</b> To use this case, the table CREATE\$JAVA\$LOB\$TABLE must exist in the current schema and must have a column LOB of type BLOB and a column NAME of type VARCHAR2.
AS <i>source_text</i>		determines a sequence of characters for a Java or SQLJ source.

---

## Examples

**Java Class Example** The following statement creates a schema object containing a Java class using the name found in a Java binary file:

```
CREATE JAVA CLASS USING BFILE (bfile_dir, 'Agent.class');
```

This example assumes the directory object `bfile_dir`, which points to the operating system directory containing the Java class `Agent.class`, already exists. In this example, the name of the class determines the name of the Java class schema object.

**Java Source Example** The following statement creates a Java source schema object:

```
CREATE JAVA SOURCE NAMED "Hello" AS
  public class Hello {
    public static String hello() {
      return "Hello World";    } };
```

**Java Resource Example** The following statement creates a Java resource schema object named `APPTEXT` from a `BFILE`:

```
CREATE JAVA RESOURCE NAMED "appText"
  USING BFILE (bfile_dir, 'textBundle.dat');
```

## CREATE LIBRARY

### Syntax



**filespec:** See "filespec" on page 7-516.

### Purpose

To create a schema object associated with an operating-system shared library. The name of this schema object can then be used in the *call\_spec* of `CREATE FUNCTION` or `CREATE PROCEDURE` statements, or when declaring a function or procedure in a package or type, so that SQL and PL/SQL can call to third-generation-language (3GL) functions and procedures.

**See Also:** "`CREATE FUNCTION`" on page 7-284 and *PL/SQL User's Guide and Reference* for more information on functions and procedures.

### Prerequisites

To create a library in your own schema, you must have the `CREATE LIBRARY` system privilege. To create a library in another user's schema, you must have the `CREATE ANY LIBRARY` system privilege. To use the procedures and functions stored in the library, you must have `EXECUTE` object privileges on the library.

The `CREATE LIBRARY` statement is valid only on platforms that support shared libraries and dynamic linking.

### Keywords and Parameters

`OR REPLACE`

re-creates the library if it already exists. Use this clause to change the definition of an existing library without dropping, re-creating, and regranted schema object privileges granted on it.

Users who had previously been granted privileges on a redefined library can still access the library without being regranted the privileges.



---

<i>libname</i>	is the name you wish to create to represent this library when declaring a function or procedure with a <i>call_spec</i> .
<i>'filespec'</i>	is a string literal, enclosed in single quotes. This string should be the path or filename your operating system recognizes as naming the shared library.  The 'filespec' is not interpreted during execution of the CREATE LIBRARY statement. The existence of the library file is not checked until an attempt is made to execute a routine from it.

---

## Examples

The following statement creates library EXT\_LIB:

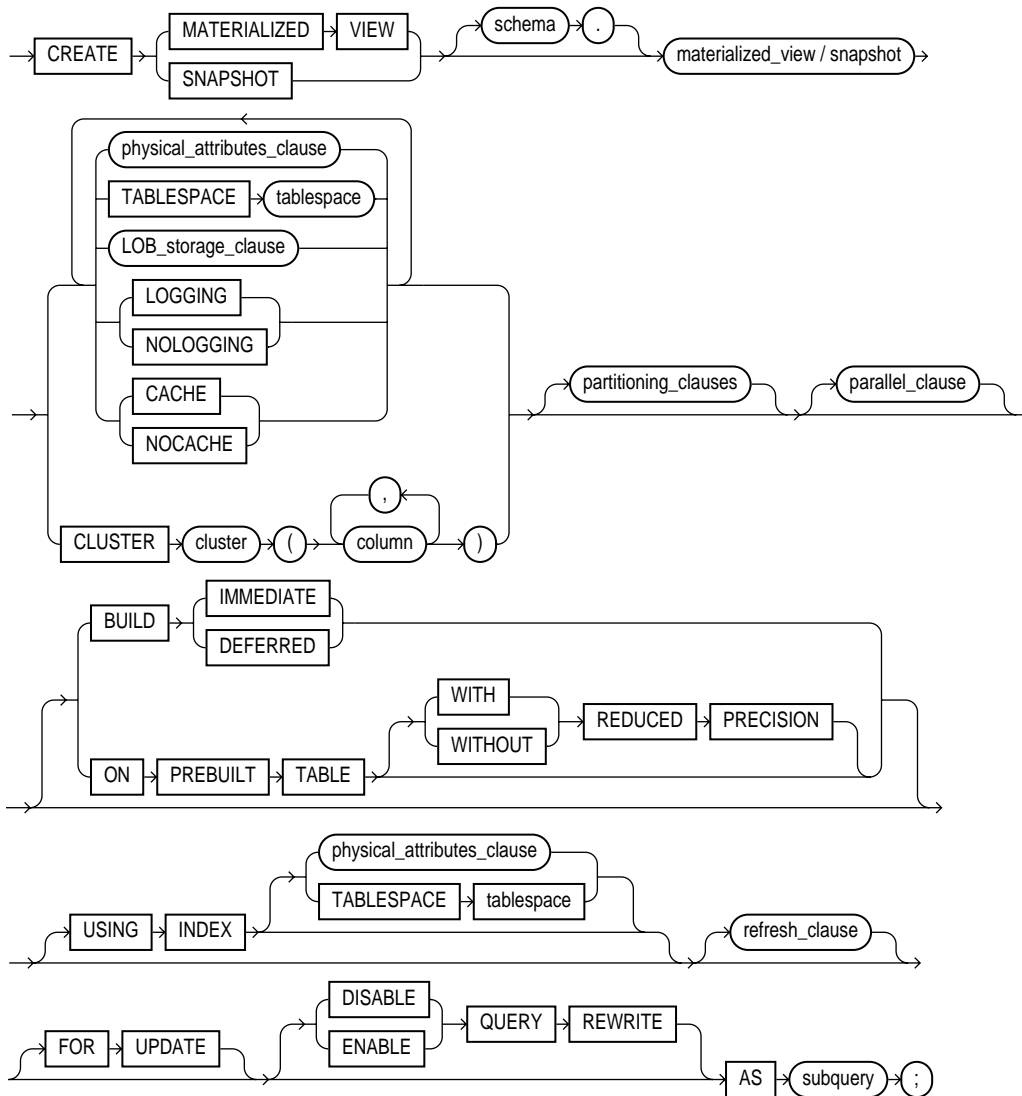
```
CREATE LIBRARY ext_lib AS '/OR/lib/ext_lib.so';
```

The following statement re-creates library EXT\_LIB:

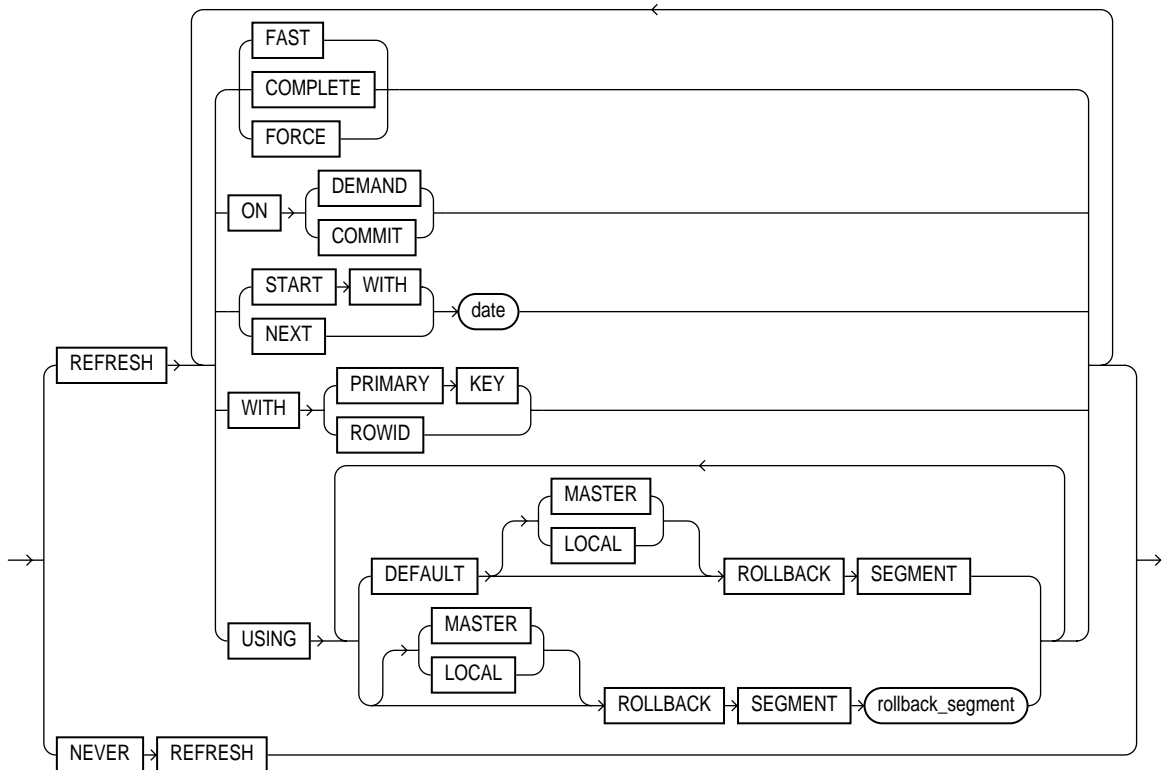
```
CREATE OR REPLACE LIBRARY ext_lib IS '/OR/newlib/ext_lib.so';
```

# CREATE MATERIALIZED VIEW / SNAPSHOT

## Syntax




**refresh\_clause::=**



**physical\_attributes\_clause:** See "CREATE TABLE" on page 4-381.

**parallel\_clause::=**



**subquery:** See "SELECT and Subqueries" on page 7-569.

**LOB\_storage\_clause:** See "CREATE TABLE" on page 4-381.

**partitioning\_clauses:** See "CREATE TABLE" on page 4-381.

## Purpose

To create a **materialized view**. The terms **snapshot** and **materialized view** are synonymous in Oracle documentation. This reference uses "materialized view" for consistency. Both refer to a database object that contains the results of a query of one or more tables.

The tables in the query are called **master tables** (a replication term) or **detail tables** (a data warehouse term). This reference uses "master tables" for consistency. The databases containing the master tables are called the **master databases**.

For replication purposes, materialized views allow you to maintain copies of remote data on your local node. The copies can be updatable with the Advanced Replication feature and are read-only without this feature. You can select data from a materialized view as you would from a table or view. In replication environments, the materialized views commonly created are **primary key**, **rowid**, and **subquery** materialized views.

For data warehousing purposes, the materialized views commonly created are **materialized aggregate views**, **single-table materialized aggregate views**, and **materialized join views**. All three types of materialized views can be used by query rewrite, an optimization technique that transforms a user request written in terms of master tables into a semantically equivalent request that includes one or more materialized view. In a data warehousing environment, all master tables must be local.

### See Also:

- *Oracle8i Replication* for information on the types of materialized views used to support replication.
- *Oracle8i Data Warehousing Guide* for information on the types of materialized views used to support data warehousing.

## Prerequisites

The privileges required to create a materialized view should be granted directly.

To create a materialized view **in your own schema**:

- You must have been granted either the `CREATE MATERIALIZED VIEW` or `CREATE SNAPSHOT` system privilege **and** either the `CREATE TABLE` or `CREATE ANY TABLE` system privilege.

- You must also have access to any master tables of the materialized view that you do not own, either through a `SELECT` object privilege on each of the tables or through the `SELECT ANY TABLE` system privilege.

To create a materialized view **in another user's schema**:

- You must have the `CREATE ANY MATERIALIZED VIEW` or `CREATE ANY SNAPSHOT` system privilege **and** access to any master tables of the materialized view that you do not own, either through a `SELECT` object privilege on each of the tables or through the `SELECT ANY TABLE` system privilege.
- The owner of the materialized view must have the `CREATE TABLE` system privilege. The owner must also have access to any master tables of the materialized view that the schema owner does not own **and** to any materialized view logs defined on those master tables, either through a `SELECT` object privilege on each of the tables or through the `SELECT ANY TABLE` system privilege.

To create the materialized view **with query rewrite enabled**, in addition to the preceding privileges:

- The owner of the master tables must have the `QUERY REWRITE` system privilege.
- If you are not the owner of the master tables, you must have the `GLOBAL QUERY REWRITE` system privilege.
- If the schema owner does not own the master tables, then the schema owner must have the `GLOBAL QUERY REWRITE` privilege.

The user whose schema contains the materialized view must have sufficient quota in the target tablespace to store the materialized view's master table and index, or must have the `UNLIMITED TABLESPACE` system privilege.

When you create a materialized view, Oracle creates one internal table and at least one index, and may create one view, all in the schema of the materialized view. Oracle uses these objects to maintain the materialized view's data. You must have the privileges necessary to create these objects.

**See Also:**

- ["CREATE TABLE"](#) on page 4-381, ["CREATE VIEW"](#) on page 7-456, and ["CREATE INDEX"](#) on page 7-291 for information on these privileges.
- *Oracle8i Replication* for information about the prerequisites that apply to creating replication materialized views.
- *Oracle8i Data Warehousing Guide* for information about the prerequisites that apply to creating data warehousing materialized views.

**Keywords and Parameters**

---

<i>schema</i>	is the schema to contain the materialized view. If you omit <i>schema</i> , Oracle creates the materialized view in your schema.
<i>materialized_view/ snapshot</i>	is the name of the materialized view to be created. Oracle generates names for the table and indexes used to maintain the materialized view by adding a prefix or suffix to the materialized view name.
<i>physical_ attributes_clause</i>	establishes values for the PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters (or, when used in the USING INDEX clause, for the INITRANS and MAXTRANS parameters only) and the storage characteristics for the materialized view.

**See Also:**

- ["CREATE TABLE"](#) on page 4-381 for information on the PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters
- ["storage\\_clause"](#) on page 7-605 for information about storage characteristics

TABLESPACE	specifies the tablespace in which the materialized view is to be created. If you omit this clause, Oracle creates the materialized view in the default tablespace of the owner of the materialized view's schema.
<i>LOB_storage_ clause</i>	specifies the LOB storage characteristics. <b>See Also:</b> <a href="#">"CREATE TABLE"</a> on page 4-381 for detailed information about specifying the parameters of this clause.
CLUSTER	creates the materialized view as part of the specified cluster. A clustered materialized view uses the cluster's space allocation. Therefore, do not use the <i>physical_attributes_clause</i> or the TABLESPACE clause with the CLUSTER clause.
LOGGING   NOLOGGING	establishes the logging characteristics for the materialized view. <b>See Also:</b> <a href="#">"CREATE TABLE"</a> on page 4-381 for a description of logging characteristics.

CACHE   NOCACHE	<p>For data that will be accessed frequently, <code>CACHE</code> specifies that the blocks retrieved for this table are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables. <code>NOCACHE</code> specifies that the blocks are placed at the least recently used end of the LRU list.</p> <p><b>See Also:</b> "<a href="#">CREATE TABLE</a>" on page 4-381 for information about specifying <code>CACHE</code> or <code>NOCACHE</code>.</p>
<i>partitioning clauses</i>	<p>specifies that the materialized view is partitioned on specified ranges of values or on a hash function. Partitioning of materialized views is the same as partitioning tables, as described in "<a href="#">CREATE TABLE</a>" on page 4-381.</p>
<i>parallel_clause</i>	<p>indicates whether parallel operations will be supported for the materialized view and sets the default degree of parallelism for queries and DML on the materialized view after creation.</p> <p><code>NOPARALLEL</code> specifies serial execution. This is the default.</p> <p><code>PARALLEL</code> causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the <code>PARALLEL_THREADS_PER_CPU</code> initialization parameter.</p> <p><code>PARALLEL integer</code> specifies the <b>degree of parallelism</b>, which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution servers. Normally Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i>.</p> <p><b>See Also:</b> The <a href="#">Notes</a> to the <i>parallel_clause</i> of "<a href="#">CREATE TABLE</a>" on page 4-381 for additional information.</p>
BUILD	<p>specifies when to populate the materialized view.</p> <p><code>IMMEDIATE</code> specifies that the materialized view is populated immediately. This is the default.</p> <p><code>DEFERRED</code> specifies that the materialized view will be populated by the next <code>REFRESH</code> operation. The first (deferred) refresh must always be a complete refresh. Until then, the materialized view has a staleness value of <code>UNUSABLE</code>, so it cannot be used for query rewrite.</p>
ON PREBUILT TABLE	<p>lets you register an existing table as a preinitialized materialized view. This is particularly useful for registering large materialized views in a data warehousing environment. The table must have the same name and be in the same schema as the resulting materialized view.</p> <p>If the materialized view is dropped, the preexisting table reverts to its identity as a table.</p> <p><b>Caution:</b> This clause assumes that the table object reflects the materialization of a subquery. Oracle Corporation strongly recommends that you ensure that this assumption is true in order to ensure that the materialized view correctly reflects the data in its master tables.</p>

**Restrictions:**

- Each column alias in *subquery* must correspond to a column in *table\_name*, and corresponding columns must have matching datatypes.
- If you specify this clause, you cannot specify a NOT NULL constraint for any column that is unmanaged (that is, not referenced in *subquery*) unless you also specify a default value for that column.

WITH REDUCED PRECISION lets you authorize the loss of precision that will result if the precision of the table or materialized view columns do not exactly match the precision returned by *subquery*.

WITHOUT REDUCED PRECISION requires that the precision of the table or materialized view columns match exactly the precision returned by *subquery*, or the create operation will fail. This is the default.

USING INDEX establishes the value of INITTRANS, MAXTRANS, and STORAGE parameters for the index Oracle uses to maintain the materialized view's data. If USING INDEX is not specified, then default values are used for the index.

**Restriction:** You cannot specify the PCTUSED or PCTFREE parameters in this clause.

*refresh\_clause* specifies the default methods, modes, and times for Oracle to refresh the materialized view. If a materialized view's master tables are modified, the data in a materialized view must be updated to make the materialized view accurately reflect the data currently in its master tables. This clause lets you schedule the times and specify the method and mode for Oracle to refresh the materialized view.

**Note:** This clause only sets the default refresh options. For instructions on actually implementing the refresh, refer to *Oracle8i Replication* and *Oracle8i Data Warehousing Guide*.

FAST specifies the incremental refresh method, which performs the refresh according to the changes that have occurred to the master tables. The changes are stored either in the materialized view log associated with the master table (for conventional DML changes) or in the direct loader log (for direct-load INSERTs).

You can create a materialized aggregate view even if you have not yet created materialized view logs for the underlying master tables. However, if you are creating any other type of materialized view, the CREATE statement will fail unless those materialized view logs already exist. (Oracle creates the direct loader log automatically when a direct-load INSERT takes place. No user intervention is needed.)



After create time, Oracle will perform the fast refresh for conventional DML only if the appropriate materialized view logs exist.

For both conventional DML changes and for direct-path loads, other conditions may restrict the eligibility of a materialized view for fast refresh.

**See Also:**

- *Oracle8i Replication* for restrictions on fast refresh in replication environments.
- *Oracle8i Data Warehousing Guide* for restrictions on fast refresh in data warehouse environments.

Materialized views are not eligible for fast refresh if the defining query contains an analytic function.

**See Also:** "[Analytic Functions](#)" on page 4-7.

COMPLETE	specifies the complete refresh method, which is implemented by executing the materialized view's defining query. If you request a complete refresh, Oracle performs a complete refresh even if a fast refresh is possible.
FORCE	specifies that when a refresh occurs, Oracle will perform a fast refresh if one is possible or a complete refresh otherwise. If you do not specify a refresh method (FAST, COMPLETE, or FORCE), FORCE is the default.
ON COMMIT	<p>specifies that a fast refresh is to occur whenever Oracle commits a transaction that operates on a master table of the materialized view.</p> <p><b>Restriction:</b> This clause is supported only for materialized join views and single-table materialized aggregate views.</p> <p><b>See Also:</b> <i>Oracle8i Replication</i> and <i>Oracle8i Data Warehousing Guide</i>.</p>
ON DEMAND	<p>specifies that the materialized view will be refreshed on demand by calling one of the three DBMS_MVIEW refresh procedures. If you omit both ON COMMIT and ON DEMAND, ON DEMAND is the default.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ <i>Oracle8i Supplied PL/SQL Packages Reference</i> for information on these procedures.</li> <li>■ <i>Oracle8i Data Warehousing Guide</i> on the types of materialized views you can create by specifying REFRESH ON DEMAND.</li> </ul>
START WITH	<p>If you specify ON COMMIT or ON DEMAND, you cannot also specify START WITH or NEXT.</p> <p>specifies a date expression for the first automatic refresh time.</p>

**NEXT** specifies a date expression for calculating the interval between automatic refreshes.

Both the `START WITH` and `NEXT` values must evaluate to a time in the future. If you omit the `START WITH` value, Oracle determines the first automatic refresh time by evaluating the `NEXT` expression with respect to the creation time of the materialized view. If you specify a `START WITH` value but omit the `NEXT` value, Oracle refreshes the materialized view only once. If you omit both the `START WITH` and `NEXT` values, or if you omit the *refresh\_clause* entirely, Oracle does not automatically refresh the materialized view.

**WITH PRIMARY KEY** specifies that a primary key materialized view is to be created. This is the default, and should be used in all cases except those described for `WITH ROWID`. Primary key materialized views allow materialized view master tables to be reorganized without affecting the materialized view's ability to continue to fast refresh. The master table must contain an enabled primary key constraint.

**See Also:** *Oracle8i Replication* for detailed information about primary key materialized views

**WITH ROWID** specifies that a rowid materialized view is to be created. Rowid materialized views provide compatibility with master tables in releases of Oracle prior to 8.0.

You can also use rowid materialized views if the materialized view does not include all primary key columns of the master tables. Rowid materialized views must be based on a single remote table and cannot contain any of the following:

- distinct or aggregate functions
- `GROUP BY` or `CONNECT BY` clauses
- subqueries
- joins
- set operations

Rowid materialized views cannot be fast refreshed after a master table reorganization until a complete refresh has been performed.

**USING ROLLBACK SEGMENT** specifies the remote rollback segment to be used during materialized view refresh, where *rollback\_segment* is the name of the rollback segment to be used.

- **DEFAULT** specifies that Oracle will choose automatically which rollback segment to use. If you specify **DEFAULT**, you cannot specify *rollback\_segment*.

**Note:** **DEFAULT** is most useful when modifying a materialized view, as described in "[ALTER MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-47.

- **MASTER** specifies the remote rollback segment to be used at the remote master site for the individual materialized view.
- **LOCAL** specifies the remote rollback segment to be used for the local refresh group that contains the materialized view.

**See Also:** *Oracle8i Replication* for information on specifying the local materialized view rollback segment using the `DBMS_REFRESH` package.

If you do not specify **MASTER** or **LOCAL**, Oracle uses **LOCAL** by default. If you do not specify *rollback\_segment*, Oracle automatically chooses the rollback segment to be used.

The master rollback segment is stored on a per-materialized-view basis and is validated during materialized view creation and refresh. If the materialized view is complex, the master rollback segment, if specified, is ignored.

**NEVER REFRESH** prevents the materialized view from being refreshed with any Oracle refresh mechanism or procedure. If you issue a **REFRESH** statement on the materialized view, Oracle returns an error.

**FOR UPDATE** allows a subquery, primary key, or rowid materialized view to be updated. When used in conjunction with Advanced Replication, these updates will be propagated to the master.

**See Also:** *Oracle8i Replication*.

**QUERY REWRITE** specifies whether the materialized view is eligible to be used for query rewrite.

**ENABLE** enables the materialized view for query rewrite.

**See Also:** *Oracle8i Data Warehousing Guide* for more information on query rewrite.

---

**Notes:**

- Query rewrite is disabled by default, so you must specify this clause to make materialized views eligible for query rewrite.
- Be sure to analyze the materialized view after you create it. Oracle needs the statistics generated by the **ANALYZE** operation to optimize query rewrite.

**Restrictions:**

- You can enable query rewrite only if all user-defined functions in the materialized view are `DETERMINISTIC`.
- You can enable query rewrite only if expressions in the statement are repeatable. For example, you cannot include `CURRENT_TIME` or `USER`, sequence values (such as the `CURRVAL` or `NEXTVAL` pseudocolumns), or the `SAMPLE` clause (which may sample different rows as the contents of the materialized view change).

**See Also:** "[CREATE FUNCTION](#)" on page 7-284 and *Oracle8i Data Warehousing Guide*.

`DISABLE` specifies that the materialized view is not eligible for use by query rewrite. However, a disabled materialized view can be refreshed.

**AS *subquery***

specifies the materialized view's defining query. When you create the materialized view, Oracle executes this query and places the results in the materialized view. This query is any valid SQL query. However, not all queries are fast refreshable, nor are all queries eligible for query rewrite.

---

**Notes:**

- Oracle does not execute the query immediately if you specify `BUILD DEFERRED`.
- Oracle recommends that you qualify each table and view in the `FROM` clause of the materialized view query with the schema containing it. For some additional caveats, see the `AS subquery` clause of "[CREATE TABLE](#)" on page 4-381.

---

**Restrictions:**

- A materialized view query can select from tables or views owned by the user `SYS`, but you cannot enable `QUERY REWRITE` on such a materialized view.
- You cannot refer to a user-defined type anywhere in the materialized view query.
- Materialized join views and materialized aggregate views with a `GROUP BY` clause cannot select from an index-organized table.
- Materialized views cannot contain columns of datatype `LONG`.
- If the subquery refers to a temporary table, you cannot create a materialized view log for this materialized view, nor can you specify the `QUERY REWRITE` clause of `CREATE MATERIALIZED VIEW` or `ALTER MATERIALIZED VIEW`.
- If the `FROM` list of the materialized view references another materialized view, you must control the refresh order of the materialized views manually. That is, you must refresh the materialized view depended upon and then the dependent materialized view in order to maintain freshness.

If you are creating a materialized view enabled for query rewrite:

- The subquery cannot contain (either directly or through a view) references to ROWNUM, USER, SYSDATE, remote tables, sequences, or PL/SQL functions that write or read database or package state.
- The materialized view and the master tables of the materialized view must be local.

If you want the materialized view to be eligible for fast refresh using a materialized view log, some additional restrictions may apply.

**See Also:**

- *Oracle8i Data Warehousing Guide* for more information on restrictions relating to data warehousing
  - *Oracle8i Replication* for more information on restrictions relating to replication.
- 

## Examples

**Materialized Aggregate View Examples** The following statement creates and populates a materialized aggregate view and specifies the default refresh method, mode, and time:

```
CREATE MATERIALIZED VIEW mv1 REFRESH FAST ON COMMIT
  BUILD IMMEDIATE
  AS SELECT t.month, p.prod_name, SUM(f.sales) AS sum_sales
     FROM time t, product p, fact f
     WHERE f.curDate = t.curDate AND f.item = p.item
     GROUP BY t.month, p.prod_name;
```

The following statement creates and populates the materialized aggregate view SALES\_BY\_MONTH\_BY\_STATE. The materialized view will be populated with data as soon as the statement executes successfully. By default, subsequent refreshes will be accomplished by reexecuting the materialized view's query:

```
CREATE MATERIALIZED VIEW sales_by_month_by_state
  TABLESPACE my_ts PARALLEL (10)
  ENABLE QUERY REWRITE
  BUILD IMMEDIATE
  REFRESH COMPLETE
  AS SELECT t.month, g.state, SUM(f.sales) AS sum_sales
     FROM fact f, time t, geog g
     WHERE f.cur_date = t.cur_date AND f.city_id = g.city_id
     GROUP BY month, state;
```

**Prebuilt Materialized View Example** The following statement creates a materialized aggregate view for the preexisting summary table, SALES\_SUM\_TABLE:

```
CREATE TABLE sales_sum_table
    (month DATE, state VARCHAR2(25), sales NUMBER);

CREATE MATERIALIZED VIEW sales_sum_table
    ON PREBUILT TABLE
    ENABLE QUERY REWRITE
    AS SELECT t.month, g.state, SUM(f.sales) AS sum_sales
        FROM fact f, time t, geog g
        WHERE f.cur_date = t.cur_date AND f.city_id = g.city_id
        GROUP BY month, state;
```

In this example, the materialized view has the same name as the prebuilt table and also has the same number of columns with the same datatypes as the prebuilt table.

**Materialized Join View Example** The following statement creates the materialized join view MJV:

```
CREATE MATERIALIZED VIEW mjv
    REFRESH FAST
    AS SELECT l.rowid as l_rid, l.pk, l.ofk, l.c1, l.c2,
        o.rowid as o_rid, o.pk, o.cfk, o.c1, o.c2,
        c.rowid as c_rid, c.pd, c.c1, c.c2
    FROM l, o, c
    WHERE l.ofk = o.pk(+) AND o.ofk = c.pk(+);
```

**Subquery Materialized View Example** The following statement creates a subquery materialized view based on the ORDERS and CUSTOMERS tables in the SALES schema at a remote database:

```
CREATE MATERIALIZED VIEW sales.orders FOR UPDATE
    AS SELECT * FROM sales.orders@dbs1.acme.com o
    WHERE EXISTS
        (SELECT * FROM sales.customers@dbs1.acme.com c
        WHERE o.c_id = c.c_id);
```

**Primary Key Example** The following statement creates the primary-key materialized view HUMAN\_GENOME:

```
CREATE MATERIALIZED VIEW human_genome
    REFRESH FAST START WITH SYSDATE NEXT SYSDATE + 1/4096
    WITH PRIMARY KEY
    AS SELECT * FROM genome_catalog;
```

**Rowid Example** The following statement creates a rowid materialized view:

```
CREATE MATERIALIZED VIEW emp_data REFRESH WITH ROWID
  AS SELECT * FROM emp_table73;
```

**Periodic Refresh Example** The following statement creates the primary key materialized view EMP\_SF and populates it with data from SCOTT's employee table in New York:

```
CREATE MATERIALIZED VIEW emp_sf
  PCTFREE 5 PCTUSED 60
  TABLESPACE users
  STORAGE (INITIAL 50K NEXT 50K)
  REFRESH FAST NEXT sysdate + 7
  AS SELECT * FROM scott.emp@ny;
```

The statement does not include a `START WITH` parameter, so Oracle determines the first automatic refresh time by evaluating the `NEXT` value using the current `SYSDATE`. Provided that a materialized view log currently exists for the employee table in New York, Oracle performs a fast refresh of the materialized view every 7 days, beginning 7 days after the materialized view is created.

Because the materialized view conforms to the conditions for fast refresh, Oracle will perform a fast refresh. The above statement also establishes storage characteristics that Oracle uses to maintain the materialized view.

**Automatic Refresh Times Example** The following statement creates the complex materialized view ALL\_EMPS that queries the employee tables in Dallas and Baltimore:

```
CREATE MATERIALIZED VIEW all_emps
  PCTFREE 5 PCTUSED 60
  TABLESPACE users
  STORAGE INITIAL 50K NEXT 50K
  USING INDEX STORAGE (INITIAL 25K NEXT 25K)
  REFRESH START WITH ROUND(SYSDATE + 1) + 11/24
  NEXT NEXT_DAY(TRUNC(SYSDATE, 'MONDAY') + 15/24
  AS SELECT * FROM fran.emp@dallas
  UNION
  SELECT * FROM marco.emp@balt;
```

Oracle automatically refreshes this materialized view tomorrow at 11:00 am and subsequently every Monday at 3:00 pm. The default refresh method is `FORCE`. `ALL_EMPS` contains a `UNION`, which is not supported for fast refresh, so Oracle will automatically perform a complete refresh.

The above statement also establishes storage characteristics for both the materialized view and the index that Oracle uses to maintain it:

- The first *storage\_clause* establishes the sizes of the first and second extents of the materialized view as 50 kilobytes each.
- The second *storage\_clause* (appearing with the USING INDEX clause) establishes the sizes of the first and second extents of the index as 25 kilobytes each.

**Rollback Segment Example** The following statement creates the primary key materialized view SALES\_EMP with rollback segment MASTER\_SEG at the remote master and rollback segment SNAP\_SEG for the local refresh group that contains the materialized view:

```
CREATE MATERIALIZED VIEW sales_emp
  REFRESH FAST START WITH SYSDATE NEXT SYSDATE + 7
  USING MASTER ROLLBACK SEGMENT master_seg
  LOCAL ROLLBACK SEGMENT snap_seg
  AS SELECT * FROM bar;
```

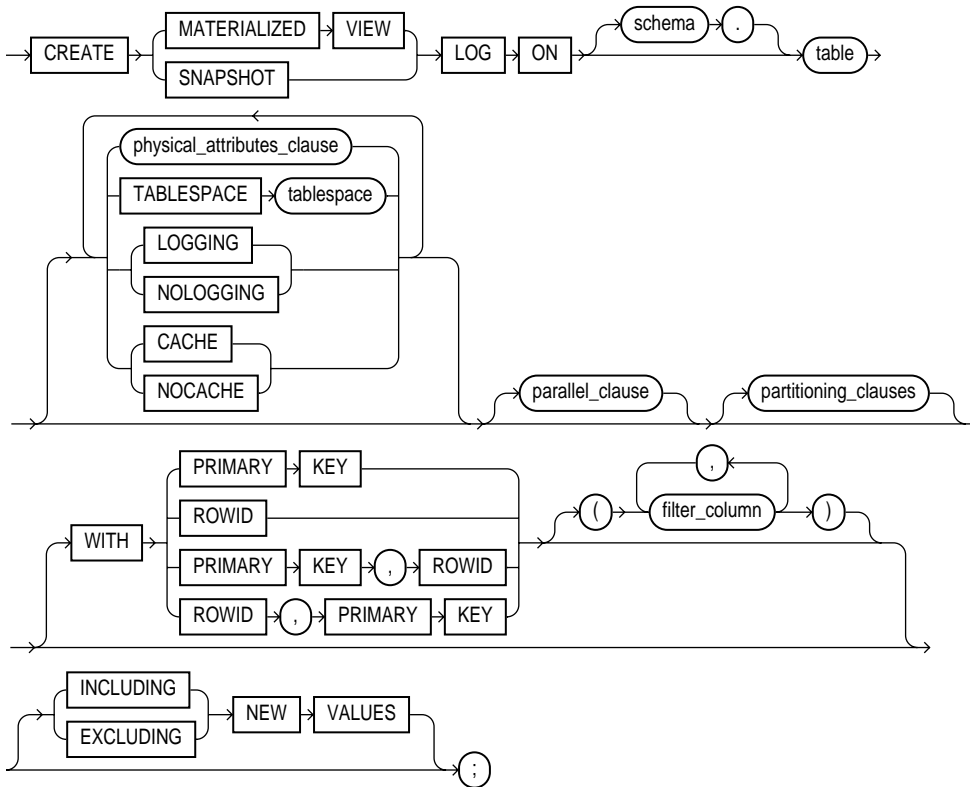
The following statement is incorrect and generates an error because it specifies a segment name with a DEFAULT rollback segment:

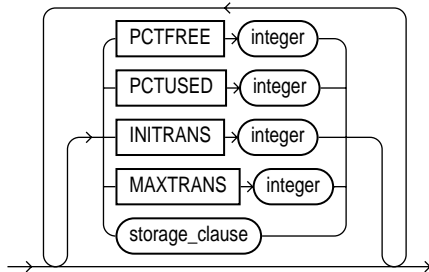
```
CREATE MATERIALIZED VIEW bogus
  REFRESH FAST START WITH SYSDATE NEXT SYSDATE + 7
  USING DEFAULT ROLLBACK SEGMENT snap_seg
  AS SELECT * FROM faux;
```




## CREATE MATERIALIZED VIEW LOG / SNAPSHOT LOG

### Syntax



**physical\_attributes\_clause::=**

**storage\_clause:** See "storage\_clause" on page 7-605.

**parallel\_clause::=**

**partitioning\_clauses:** See "CREATE TABLE" on page 4-381.

**Purpose**

To create a materialized view log. A **materialized view log** is a table associated with the master table of a materialized view. The terms **snapshot** and **materialized view** are synonymous. Both refer to a table that contains the results of a query of one or more tables, each of which may be located on the same or on a remote database.

When DML changes are made to the master table's data, Oracle stores rows describing those changes in the materialized view log and then uses the materialized view log to refresh materialized views based on the master table. This process is called a *fast refresh*. Without a materialized view log, Oracle must reexecute the materialized view query to refresh the materialized view. This process is called a *complete refresh*. Usually, a fast refresh takes less time than a complete refresh.

A materialized view log is located in the master database in the same schema as the master table. You need only a single materialized view log for a master table. Oracle can use this materialized view log to perform fast refreshes for all fast-refreshable materialized views based on the master table.

To fast refresh a materialized join view (a materialized view containing a join), you must create a materialized view log for each of its master tables.

**See Also:**

- *Oracle8i Data Warehousing Guide* and *Oracle8i Replication* for information on materialized views in general.
- "[CREATE MATERIALIZED VIEW / SNAPSHOT](#)" on page 7-318 and *Oracle8i Concepts* for more information on materialized views.
- "[ALTER MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-58 for information on modifying a materialized view log.
- "[DROP MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-487 for information on dropping a materialized view log.
- *Oracle8i Concepts* for information on using direct loader logs.

## Prerequisites

The privileges required to create a materialized view log directly relate to the privileges necessary to create the underlying objects associated with a materialized view log.

- If you own the master table, you can create an associated materialized view log if you have the `CREATE TABLE` privilege.
- If you are creating a materialized view log for a table in another user's schema, you must have the `CREATE ANY TABLE` and `COMMENT ANY TABLE` privileges, as well as either the `SELECT` privilege for the master table or `SELECT ANY TABLE`.

In either case, the owner of the materialized view log must have sufficient quota in the tablespace intended to hold the materialized view log or must have the `UNLIMITED TABLESPACE` system privilege.

For detailed information about the prerequisites for creating a materialized view log, see *Oracle8i Replication*.

## Keywords and Parameters

<i>schema</i>	is the schema containing the materialized view log's master table. If you omit <i>schema</i> , Oracle assumes the master table is contained in your own schema. Oracle creates the materialized view log in the schema of its master table. You cannot create a materialized view log for a table in the schema of the user <i>SYS</i> .
<i>table</i>	is the name of the master table for which the materialized view log is to be created. You cannot create a materialized view log for a view.
<i>physical_</i> <i>attributes_clause</i>	establishes values for physical and storage characteristics for the materialized view log. See the descriptions of these parameters in " <a href="#">CREATE TABLE</a> " on page 4-381 and " <a href="#">storage_clause</a> " on page 7-605.
TABLESPACE	specifies the tablespace in which the materialized view log is to be created. If you omit this clause, Oracle creates the materialized view log in the default tablespace of the owner of the materialized view log's schema.
LOGGING   NOLOGGING	establishes the logging characteristics for the materialized view log. For a description of logging characteristics, see " <a href="#">CREATE TABLE</a> " on page 4-381.
CACHE   NOCACHE	For data that will be accessed frequently, <i>CACHE</i> specifies that the blocks retrieved for this log are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This attribute is useful for small lookup tables. <i>NOCACHE</i> specifies that the blocks are placed at the least recently used end of the LRU list. <b>See Also:</b> " <a href="#">ALTER TABLE</a> " on page 7-123 for information about specifying <i>CACHE</i> or <i>NOCACHE</i> .
<i>parallel_clause</i>	specifies whether parallel operations will be supported for the materialized view log. For additional information, see the <a href="#">Notes</a> to the <i>parallel_clause</i> of " <a href="#">CREATE TABLE</a> " on page 4-381.
NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the <i>PARALLEL_THREADS_PER_CPU</i> initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution servers. Normally, Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .
<i>partitioning_</i> <i>clauses</i>	specifies that the materialized view log is partitioned on specified ranges of values or on a hash function. Partitioning of materialized view logs is the same as partitioning of tables, as described in " <a href="#">CREATE TABLE</a> " on page 4-381.
WITH	specifies whether the materialized view log should record the primary key, the rowid, or both the primary key and rowid when rows in the master are updated.

---

This clause also specifies whether the materialized view log records filter columns, which are non-primary-key columns referenced by subquery materialized views.

**PRIMARY KEY** specifies that the primary key of all rows updated should be recorded in the materialized view log. The primary key of updated rows in the master table must be recorded in the materialized view log.

**ROWID** specifies that the rowid of all rows updated should be recorded in the materialized view log. The rowid must be recorded in the materialized view log.

*filter\_column* is a comma-separated list that specifies the filter columns to be recorded in the materialized view log. For fast-refreshable primary-key materialized views defined with subqueries, all filter columns referenced by the defining subquery must be recorded in the materialized view log.

Oracle records the primary key of all rows updated in the master by default.

**NEW VALUES** specifies whether Oracle saves both old and new values in the materialized view log.

**INCLUDING** saves both new and old values in the log. If this log is for a table on which you have a single-table materialized aggregate view, and if you want the materialized view to be eligible for fast refresh, you must specify **INCLUDING**.

**EXCLUDING** disables the recording of new values in the log. This is the default. You can use this clause to avoid the overhead of recording new values. However, do not use this clause if you have a fast-refreshable single-table materialized aggregate view defined on this table.

---

## Examples

**Primary Key Examples** The following statement creates a materialized view log on an employee table that records only primary key values:

```
CREATE MATERIALIZED VIEW LOG ON emp WITH PRIMARY KEY;
```

Oracle can use this materialized view log to perform a fast refresh on any simple primary key materialized view subsequently created on the **EMP** table.

The following statement also creates a materialized view log that record only the primary keys of updated rows:

```
CREATE MATERIALIZED VIEW LOG ON emp
  PCTFREE 5
  TABLESPACE users
  STORAGE (INITIAL 10K NEXT 10K);
```

**ROWID Example** The following statement creates a materialized view log that records both the primary keys and the rowids of updated rows:

```
CREATE MATERIALIZED VIEW LOG ON sales WITH ROWID, PRIMARY KEY;
```

**Filter Column Example** The following statement creates a materialized view log that records primary keys and updates to the filter column ZIP:

```
CREATE MATERIALIZED VIEW LOG ON address WITH (zip);
```

**NEW VALUES Example** The following example creates a master table, then creates a materialized view log that specifies INCLUDING NEW VALUES:

```
CREATE TABLE agg
  (u NUMBER, a NUMBER, b NUMBER, c NUMBER, d NUMBER);

CREATE MATERIALIZED VIEW LOG ON agg
  WITH ROWID (u,a,b,c,d)
  INCLUDING NEW VALUES;
```


You could create the following materialized aggregate view to use the AGG log:

```
CREATE MATERIALIZED VIEW sn0
  REFRESH FAST ON COMMIT
  AS SELECT SUM(b+c), COUNT(*), a, d, COUNT(b+c)
     FROM agg
     GROUP BY a,d;
```


This materialized view is eligible for fast refresh because the log it uses includes both old and new values.

# CREATE OPERATOR


## Syntax




**binding\_clause::=**




**implementation\_clause::=**



**context\_clause::=**



**using\_clause::=**



## Purpose

To create a new operator and define its bindings.

Operators can be referenced by indextypes and by DML and query SQL statements. The operators, in turn, reference functions, packages, types, and other user-defined objects. For a discussion of these dependencies, and of operators in general, see *Oracle8i Data Cartridge Developer's Guide* and *Oracle8i Concepts*.

## Prerequisites

To create an operator in your own schema, you must have `CREATE OPERATOR` system privilege. To create an operator in another schema, you must have the `CREATE ANY OPERATOR` system privilege. In either case, you must also have `EXECUTE` privilege on the functions and operators referenced.

## Keywords and Parameters

---

<code>OR REPLACE</code>	replaces the definition of the operator schema object. <b>Restriction:</b> You can replace the definition only if the operator has no dependent objects (for example, indextypes supporting the operator).
<i>schema</i>	is the schema containing the operator. If you omit <i>schema</i> , Oracle assumes the operator is in your own schema.
<i>operator</i>	is the name of the operator to be created.
<i>binding_clause</i>	specifies one or more parameter datatypes ( <i>parameter_type</i> ) for binding the operator to a function. The signature of each binding (that is, the sequence of the datatypes of the arguments to the corresponding function) must be unique according to the rules of overloading.  The <i>parameter_type</i> can itself be an object type. If it is, you can optionally qualify it with its schema. <b>Restriction:</b> You cannot specify a <i>parameter_type</i> of <code>REF</code> , <code>LONG</code> , or <code>LONG RAW</code> . <b>See Also:</b> <i>PL/SQL User's Guide and Reference</i> for more information about overloading.
<code>RETURN</code>	specifies the return datatype ( <i>return_type</i> ) for the binding.  The <i>return_type</i> can itself be an object type. If so, you can optionally qualify it with its schema. <b>Restriction:</b> You cannot specify a <i>return_type</i> of <code>REF</code> , <code>LONG</code> , or <code>LONG RAW</code> .



---

*implementation\_clause*

*ANCILLARY TO primary\_operator* specifies that the operator binding is ancillary to the specified primary operator binding (*primary\_operator*). If you specify this clause, do not specify a previous binding with just one number parameter.

*context\_clause* specifies the name of the implementation type used by the functional implementation of the operator as a scan context.

COMPUTE  
ANCILLARY  
DATA specifies that the operator binding computes ancillary data.

*using\_clause* specifies the function that provides the implementation for the binding.

*function\_name* is the name of the function. The function can be a standalone function, packaged function, type method, or a synonym for any of these.

---


## Example

This example creates an operator called `MERGE` in the `SCOTT` schema with two bindings. The first binding is for merging two `VARCHAR2` values and returning a `VARCHAR2` result. The second binding is for merging two geometries into a single geometry. The corresponding functional implementations for the bindings are also specified.

```
CREATE OPERATOR scott.merge
BINDING (varchar2, varchar2) RETURN varchar2
        USING text.merge,
        (spatial.geo, spatial.geo) RETURN spatial.geo
        USING spatial.merge;
```

## CREATE OUTLINE

### Syntax



### Purpose

To create a **stored outline**, which is a set of attributes used by the optimizer to generate an execution plan. You can then instruct the optimizer to use a set of outlines to influence the generation of execution plans whenever a particular SQL statement is issued, regardless of changes in factors that can affect optimization. (To modify an outline so that it takes into account changes in these factors, see "[ALTER OUTLINE](#)" on page 7-63.)

You enable or disable the use of stored outlines dynamically for an individual session or for the system. See "[ALTER SESSION](#)" on page 7-83 and "[ALTER SYSTEM](#)" on page 7-102.

**See Also:** *Oracle8i Designing and Tuning for Performance.*

### Prerequisites

To create an outline, you must have the `CREATE ANY OUTLINE` system privilege.

### Keywords and Parameters

<code>OR REPLACE</code>	replaces an existing outline with a new outline of the same name.
<i>outline</i>	is the unique name to be assigned to the stored outline. If you do not specify <i>outline</i> , the system generates an outline name.
<code>FOR CATEGORY</code> <i>category</i>	specifies an optional name used to group stored outlines. For example, you could specify a category of outlines for end-of-week use and another for end-of-quarter use. If you do not specify <i>category</i> , the outline is stored in the <code>DEFAULT</code> category.
<code>ON statement</code>	is the SQL statement for which Oracle will create an outline when the statement is compiled. You can specify any one of the following statements:

- 
- SELECT
  - DELETE
  - UPDATE
  - INSERT ... SELECT
  - CREATE TABLE ... AS SELECT
- 

**Note:** You can specify multiple outlines for a single statement, but each outline for the same statement must be in a different category.

---

## Example


The following statement creates a stored outline by compiling the `ON` statement. The outline is called `SALARIES` and is stored in the category `SPECIAL`.

```
CREATE OUTLINE salaries FOR CATEGORY special
  ON SELECT ename, sal FROM emp;
```


When this same `SELECT` statement is subsequently compiled, if the `USE_STORED_OUTLINES` parameter is set to `SPECIAL`, Oracle generates the same execution plan as was generated when the outline `SALARIES` was created.

## CREATE PACKAGE

### Syntax



### invoker\_rights\_clause::=



### Purpose

To create the specification for a stored package. A **package** is an encapsulated collection of related procedures, functions, and other program objects stored together in the database. The **specification** declares these objects.

For information on creating standalone functions and procedures, see "[CREATE FUNCTION](#)" on page 7-284 and "[CREATE PROCEDURE](#)" on page 7-353. For information on modifying a package, see "[ALTER PACKAGE](#)" on page 7-64. For information on dropping a package, see "[DROP PACKAGE](#)" on page 7-491.

For detailed discussions of packages and how to use them, see *Oracle8i Application Developer's Guide - Fundamentals* and *Oracle8i Supplied PL/SQL Packages Reference*.

### Prerequisites

Before a package can be created, the user SYS must run the SQL script DBMSSTD<sub>X</sub>.SQL. The exact name and location of this script depend on your operating system.

To create a package in your own schema, you must have CREATE PROCEDURE system privilege. To create a package in another user's schema, you must have CREATE ANY PROCEDURE system privilege.

To embed a `CREATE PACKAGE` statement inside an Oracle precompiler program, you must terminate the statement with the keyword `END-EXEC` followed by the embedded SQL statement terminator for the specific language.

**See Also:** *PL/SQL User's Guide and Reference.*

## Keywords and Parameters

<code>OR REPLACE</code>	<p>re-creates the package specification if it already exists. Use this clause to change the specification of an existing package without dropping, re-creating, and regranting object privileges previously granted on the package. If you change a package specification, Oracle recompiles it. For information on recompiling package specifications, see "<a href="#">ALTER PACKAGE</a>" on page 7-64.</p> <p>Users who had previously been granted privileges on a redefined package can still access the package without being regranted the privileges.</p> <p>If any function-based indexes depend on the package, Oracle marks the indexes <code>DISABLED</code>.</p>
<i>schema</i>	<p>is the schema to contain the package. If you omit <i>schema</i>, Oracle creates the package in your own schema.</p>
<i>package</i>	<p>is the name of the package to be created.</p> <p>If creating the package results in compilation errors, Oracle returns an error. You can see the associated compiler error messages with the <code>SHOW ERRORS</code> command.</p>
<i>invoker_rights_clause</i>	<p>lets you specify whether the functions and procedures in the package execute with the privileges and in the schema of the user who owns it or with the privileges and in the schema of <code>CURRENT_USER</code>. This specification applies to the corresponding package body as well. (For information on how <code>CURRENT_USER</code> is determined, see <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i>.)</p> <p>This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the package.</p> <p><b>See Also:</b> <i>PL/SQL User's Guide and Reference.</i></p>
<code>AUTHID CURRENT_USER</code>	<p>specifies that the package executes with the privileges of <code>CURRENT_USER</code>. This clause creates an "invoker-rights package."</p> <p>This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of <code>CURRENT_USER</code>. External names in all other statements resolve in the schema in which the package resides.</p>
<code>AUTHID DEFINER</code>	<p>specifies that the package executes with the privileges of the owner of the schema in which the package resides and that external names resolve in the schema where the package resides. This is the default.</p>

*pl/sql\_package\_*  
*spec*

is the package specification, which can contain type definitions, cursor declarations, variable declarations, constant declarations, exception declarations, PL/SQL subprogram specifications, and call specifications (declarations of a C or Java routine expressed in PL/SQL).

For a list of restrictions on user-defined functions in a package, see "[Restrictions on User-Defined Functions](#)" on page 7-286.

**See Also:**

- *PL/SQL User's Guide and Reference* for more information on PL/SQL package program units.
  - *Oracle8i Supplied PL/SQL Packages Reference* for information on Oracle supplied packages.
- 

## Example

The following SQL statement creates the specification of the EMP\_MGMT package:

```
CREATE PACKAGE emp_mgmt AS
    FUNCTION hire(ename VARCHAR2, job VARCHAR2, mgr NUMBER,
                 sal NUMBER, comm NUMBER, deptno NUMBER)
        RETURN NUMBER;
    FUNCTION create_dept(dname VARCHAR2, loc VARCHAR2)
        RETURN NUMBER;
    PROCEDURE remove_emp(empno NUMBER);
    PROCEDURE remove_dept(deptno NUMBER);
    PROCEDURE increase_sal(empno NUMBER, sal_incr NUMBER);
    PROCEDURE increase_comm(empno NUMBER, comm_incr NUMBER);
        no_comm EXCEPTION;
        no_sal EXCEPTION;
END emp_mgmt;
```

The specification for the EMP\_MGMT package declares the following public program objects:

- the functions HIRE and CREATE\_DEPT
- the procedures REMOVE\_EMP, REMOVE\_DEPT, INCREASE\_SAL, and INCREASE\_COMM
- the exceptions NO\_COMM and NO\_SAL


All of these objects are available to users who have access to the package. After creating the package, you can develop applications that call any of the package's public procedures or functions or raise any of the package's public exceptions.

Before you can call this package's procedures and functions, you must define these procedures and functions in the package body. For an example of a CREATE

PACKAGE BODY statement that creates the body of the EMP\_MGMT package, see ["CREATE PACKAGE BODY"](#) on page 7-348.

## CREATE PACKAGE BODY

### Syntax



### Purpose

To create the body of a stored package. A **package** is an encapsulated collection of related procedures, stored functions, and other program objects stored together in the database. The **body** defines these objects. For information on creating standalone functions and procedures, see ["CREATE FUNCTION"](#) on page 7-284 and ["CREATE PROCEDURE"](#) on page 7-353.

Packages are an alternative to creating procedures and functions as standalone schema objects. For a discussion of packages, including how to create packages, see ["CREATE PACKAGE"](#) on page 7-344. For some illustrations, see ["Examples"](#) on page 7-349.

For information on modifying a package, see ["ALTER PACKAGE"](#) on page 7-64. For information on removing a package from the database, see ["DROP PACKAGE"](#) on page 7-491.

### Prerequisites

Before a package can be created, the user SYS must run the SQL script DBMSSTD<sub>X</sub>.SQL. The exact name and location of this script depend on your operating system.

To create a package in your own schema, you must have CREATE PROCEDURE system privilege. To create a package in another user's schema, you must have CREATE ANY PROCEDURE system privilege.

To embed a CREATE PACKAGE BODY statement inside an Oracle precompiler program, you must terminate the statement with the keyword END-EXEC followed by the embedded SQL statement terminator for the specific language.



**See Also:** *PL/SQL User's Guide and Reference.*

## Keywords and Parameters

---

<code>OR REPLACE</code>	<p>re-creates the package body if it already exists. Use this clause to change the body of an existing package without dropping, re-creating, and regranted object privileges previously granted on it. If you change a package body, Oracle recompiles it. For information on recompiling package bodies, see "<a href="#">ALTER PACKAGE</a>" on page 7-64.</p> <p>Users who had previously been granted privileges on a redefined package can still access the package without being regranted the privileges.</p>
<code>schema</code>	is the schema to contain the package. If you omit <code>schema</code> , Oracle creates the package in your current schema.
<code>package</code>	is the name of the package to be created.
<code>pl/sql_package_body</code>	<p>is the package body, which can contain PL/SQL subprogram bodies or call specifications (declarations of a C or Java routine expressed in PL/SQL).</p> <p>For a list of restrictions on user-defined functions in a package, see "<a href="#">Restrictions on User-Defined Functions</a>" on page 7-286.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ <i>Oracle8i Application Developer's Guide - Fundamentals</i> for more information on writing PL/SQL or C package program units.</li> <li>■ <i>Oracle8i Java Stored Procedures Developer's Guide</i> for information on JAVA package program units.</li> </ul>

---

## Examples

This SQL statement creates the body of the EMP\_MGMT package:

```
CREATE PACKAGE BODY emp_mgmt AS
    tot_emps NUMBER;
    tot_depts NUMBER;

    FUNCTION hire
        (ename VARCHAR2,
         job VARCHAR2,
         mgr NUMBER,
         sal NUMBER,
         comm NUMBER,
         deptno NUMBER)

    RETURN NUMBER IS
        new_empno NUMBER(4);
```

## CREATE PACKAGE BODY

---

```
BEGIN
    SELECT empseq.NEXTVAL
        INTO new_empno
        FROM DUAL;
    INSERT INTO emp
        VALUES (new_empno, ename, job, mgr, sal, comm, deptno,
            tot_emps := tot_emps + 1;
    RETURN(new_empno);
END;

FUNCTION create_dept(dname VARCHAR2, loc VARCHAR2)
    RETURN NUMBER IS
    new_deptno NUMBER(4);
BEGIN
    SELECT deptseq.NEXTVAL
        INTO new_deptno
        FROM dual;
    INSERT INTO dept
        VALUES (new_deptno, dname, loc);
        tot_depts := tot_depts + 1;
    RETURN(new_deptno);
END;

PROCEDURE remove_emp(empno NUMBER) IS
    BEGIN
        DELETE FROM emp
        WHERE emp.empno = remove_emp.empno;
        tot_emps := tot_emps - 1;
    END;

PROCEDURE remove_dept(deptno NUMBER) IS
    BEGIN
        DELETE FROM dept
        WHERE dept.deptno = remove_dept.deptno;
        tot_depts := tot_depts - 1;
        SELECT COUNT(*)
            INTO tot_emps
            FROM emp;
        /* In case Oracle deleted employees from the EMP table
        to enforce referential integrity constraints, reset
        the value of the variable TOT_EMPS to the total
        number of employees in the EMP table. */
    END;

PROCEDURE increase_sal(empno NUMBER, sal_incr NUMBER) IS
```

```
curr_sal NUMBER(7,2);
BEGIN
    SELECT sal
    INTO curr_sal
    FROM emp
    WHERE emp.empno = increase_sal.empno;
    IF curr_sal IS NULL
        THEN RAISE no_sal;
    ELSE
        UPDATE emp
        SET sal = sal + sal_incr
        WHERE empno = empno;
    END IF;
END;

PROCEDURE increase_comm(empno NUMBER, comm_incr NUMBER) IS
    curr_comm NUMBER(7,2);
BEGIN
    SELECT comm
    INTO curr_comm
    FROM emp
    WHERE emp.empno = increase_comm.empno
    IF curr_comm IS NULL
        THEN RAISE no_comm;
    ELSE
        UPDATE emp
        SET comm = comm + comm_incr;
    END IF;
END;

END emp_mgmt;
```

This package body corresponds to the package specification in the example of the ["CREATE PACKAGE"](#) statement earlier in this chapter. The package body defines the public program objects declared in the package specification:

- the functions HIRE and CREATE\_DEPT
- the procedures REMOVE\_EMP, REMOVE\_DEPT, INCREASE\_SAL, and INCREASE\_COMM

These objects are *declared* in the package specification, so they can be called by application programs, procedures, and functions outside the package. For example, if you have access to the package, you can create a procedure INCREASE\_ALL\_

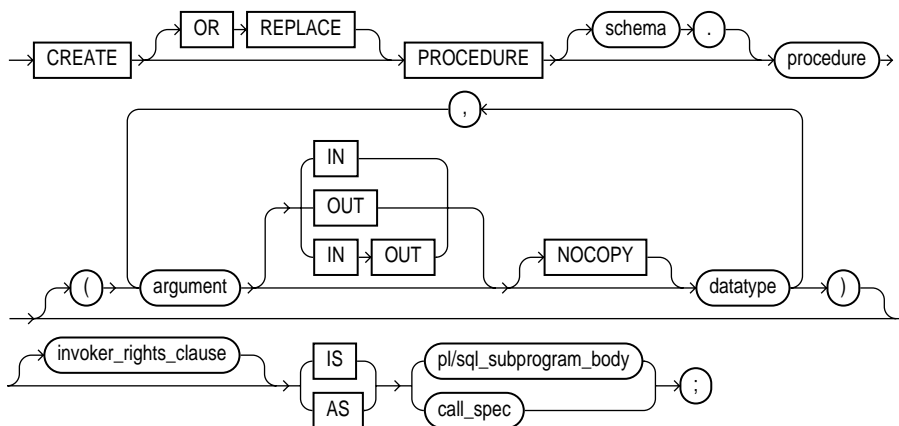
COMMS separate from the EMP\_MGMT package that calls the INCREASE\_COMM procedure.

These objects are *defined* in the package body, so you can change their definitions without causing Oracle to invalidate dependent schema objects. For example, if you subsequently change the definition of HIRE, Oracle need not recompile INCREASE\_ALL\_COMMS before executing it.


The package body in this example also declares private program objects, the variables TOT\_EMPS and TOT\_DEPTS. These objects are declared in the package body rather than the package specification, so they are accessible to other objects in the package, but they are not accessible outside the package. For example, you cannot develop an application that explicitly changes the value of the variable TOT\_DEPTS. However, the function CREATE\_DEPT is part of the package, so CREATE\_DEPT can change the value of TOT\_DEPTS.

# CREATE PROCEDURE


## Syntax




### invoker\_rights\_clause::=




### call\_spec::=



### Java\_declaration::=



**C\_declaration::=****Purpose**

To create a standalone stored procedure or a call specification.

A **procedure** is a group of PL/SQL statements that you can call by name. A **call specification** ("**call spec**") declares a Java method or a third-generation language (3GL) routine so that it can be called from SQL and PL/SQL. The call spec tells Oracle which Java method to invoke when a call is made. It also tells Oracle what type conversions to make for the arguments and return value.

Stored procedures offer advantages in the areas of development, integrity, security, performance, and memory allocation.

**See Also:**

- *Oracle8i Application Developer's Guide - Fundamentals* for more information on stored procedures, including how to call stored procedures.
- ["CREATE FUNCTION"](#) on page 7-284 for information specific to functions, which are similar in many ways.
- ["CREATE PACKAGE"](#) on page 7-344 for information on creating packages. (The `CREATE PROCEDURE` statement creates a procedure as a standalone schema object. You can also create a procedure as part of a package.)
- ["ALTER PROCEDURE"](#) on page 7-67 and ["DROP PROCEDURE"](#) on page 7-493 for information on modifying and dropping a standalone procedure.
- ["CREATE LIBRARY"](#) on page 7-316 for more information about shared libraries.
- *Oracle8i Application Developer's Guide - Fundamentals* for more information about registering external procedures.

**Prerequisites**

Before creating a procedure, the user `SYS` must run the SQL script `DBMSSTDX.SQL`. The exact name and location of this script depends on your operating system.

To create a procedure in your own schema, you must have the `CREATE PROCEDURE` system privilege. To create a procedure in another user's schema, you must have `CREATE ANY PROCEDURE` system privilege. To replace a procedure in another schema, you must have the `ALTER ANY PROCEDURE` system privilege.

To invoke a call spec, you may need additional privileges (for example, `EXECUTE` privileges on the C library for a C call spec).

To embed a `CREATE PROCEDURE` statement inside an Oracle precompiler program, you must terminate the statement with the keyword `END-EXEC` followed by the embedded SQL statement terminator for the specific language.

**See Also:** *PL/SQL User's Guide and Reference* or *Oracle8i Java Stored Procedures Developer's Guide* for more information on such prerequisites.

## Keywords and Parameters

---

<code>OR REPLACE</code>	<p>re-creates the procedure if it already exists. Use this clause to change the definition of an existing procedure without dropping, re-creating, and regrating object privileges previously granted on it. If you redefine a procedure, Oracle recompiles it. For information on recompiling procedures, see "<a href="#">ALTER PROCEDURE</a>" on page 7-67.</p> <p>Users who had previously been granted privileges on a redefined procedure can still access the procedure without being regranted the privileges.</p> <p>If any function-based indexes depend on the package, Oracle marks the indexes <code>DISABLED</code>.</p>
<i>schema</i>	is the schema to contain the procedure. If you omit <i>schema</i> , Oracle creates the procedure in your current schema.
<i>procedure</i>	is the name of the procedure to be created.
<i>argument</i>	is the name of an argument to the procedure. If the procedure does not accept arguments, you can omit the parentheses following the procedure name.
<code>IN</code>	specifies that you must specify a value for the argument when calling the procedure.
<code>OUT</code>	specifies that the procedure passes a value for this argument back to its calling environment after execution.
<code>IN OUT</code>	specifies that you must specify a value for the argument when calling the procedure and that the procedure passes a value back to its calling environment after execution.
<code>NOCOPY</code>	<p>If you omit <code>IN</code>, <code>OUT</code>, and <code>IN OUT</code>, the argument defaults to <code>IN</code>.</p> <p>instructs Oracle to pass this argument as fast as possible. This clause can significantly enhance performance when passing a large value like a record, an index-by table, or a varray to an <code>OUT</code> or <code>IN OUT</code> parameter. (<code>IN</code> parameter values are always passed <code>NOCOPY</code>.)</p> <ul style="list-style-type: none"><li>■ When you specify <code>NOCOPY</code>, assignments made to a package variable may show immediately in this parameter (or assignments made to this parameter may show immediately in a package variable) if the package variable is passed as the actual assignment corresponding to this parameter.</li><li>■ Similarly, changes made either to this parameter or to another parameter may be visible immediately through both names if the same variable is passed to both.</li><li>■ If the procedure is exited with an unhandled exception, any assignment made to this parameter may be visible in the caller's variable.</li></ul>



---

	These effects may or may not occur on any particular call. You should use <code>NOCOPY</code> only when these effects would not matter.
<i>datatype</i>	<p>is the datatype of the argument. An argument can have any datatype supported by PL/SQL.</p> <p>Datatypes cannot specify length, precision, or scale. For example, <code>VARCHAR2(10)</code> is not valid, but <code>VARCHAR2</code> is valid. Oracle derives the length, precision, and scale of an argument from the environment from which the procedure is called.</p>
<i>invoker_rights_clause</i>	<p>lets you specify whether the procedure executes with the privileges and in the schema of the user who owns it or with the privileges and in the schema of <code>CURRENT_USER</code>. (For information on how <code>CURRENT_USER</code> is determined, see <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i>.)</p> <p>This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the procedure.</p> <p><b>See Also:</b> <i>PL/SQL User's Guide and Reference</i>.</p>
	<p><code>AUTHID CURRENT_USER</code> specifies that the procedure executes with the privileges of <code>CURRENT_USER</code>. This clause creates an "invoker-rights procedure."</p> <p>This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of <code>CURRENT_USER</code>. External names in all other statements resolve in the schema in which the procedure resides.</p>
	<p><code>AUTHID DEFINER</code> specifies that the procedure executes with the privileges of the owner of the schema in which the procedure resides, and that external names resolve in the schema where the procedure resides. This is the default.</p>
<i>pl/sql_subprogram_body</i>	<p>declares the procedure in a PL/SQL subprogram body.</p> <p><b>See Also:</b> <i>Oracle8i Application Developer's Guide - Fundamentals</i> for more information on PL/SQL subprograms.</p>
<i>call_spec</i>	<p>maps a Java or C method name, parameter types, and return type to their SQL counterparts.</p> <ul style="list-style-type: none"> <li>■ In <i>Java_declaration</i>, 'string' identifies the Java implementation of the method.</li> <li>■ For an explanation of the parameters and semantics of the <i>C_declaration</i>, see <i>Oracle8i Application Developer's Guide - Fundamentals</i>.</li> </ul> <p><b>See Also:</b> <i>Oracle8i Java Stored Procedures Developer's Guide</i>.</p>
	<p><code>AS EXTERNAL</code> is an alternative way of declaring a C method. This clause has been deprecated and is supported for backward compatibility only. Oracle Corporation recommends that you use the <code>AS LANGUAGE C</code> syntax.</p>

---

## Examples

The following statement creates the procedure `CREDIT` in the schema `SAM`:

```
CREATE PROCEDURE sam.credit (acc_no IN NUMBER, amount IN NUMBER) AS
```

## CREATE PROCEDURE

---

```
BEGIN
  UPDATE accounts
  SET balance = balance + amount
  WHERE account_id = acc_no;
END;
```

The CREDIT procedure credits a specified bank account with a specified amount. When you call the procedure, you must specify the following arguments:

ACC_NO	is the number of the bank account to be credited. The argument's datatype is NUMBER.
AMOUNT	is the amount of the credit. The argument's datatype is NUMBER.


The procedure uses an UPDATE statement to increase the value in the BALANCE column of the ACCOUNTS table by the value of the argument AMOUNT for the account identified by the argument ACC\_NO.

In the following example, external procedure C\_FIND\_ROOT expects a pointer as a parameter. Procedure FIND\_ROOT passes the parameter by reference using the BY REF phrase:

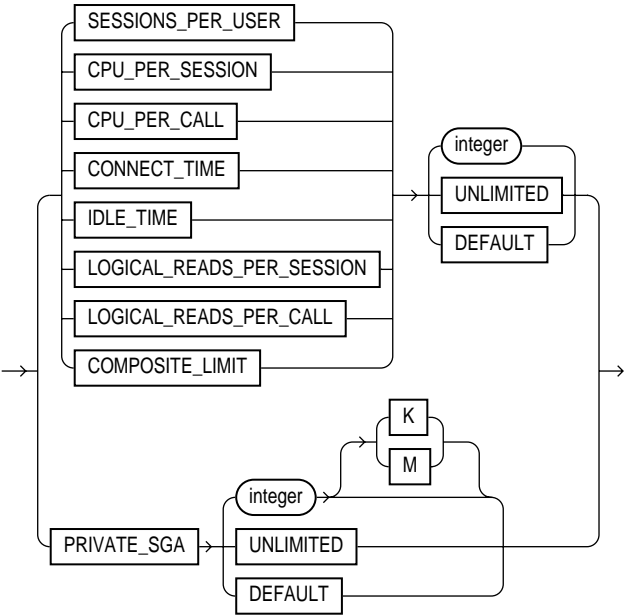
```
CREATE PROCEDURE find_root
  ( x IN REAL )
  IS LANGUAGE C
  NAME "c_find_root"
  LIBRARY c_utils
  PARAMETERS ( x BY REF );
```

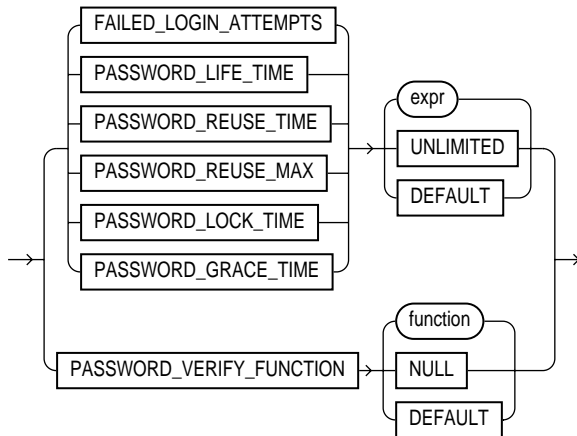
# CREATE PROFILE

## Syntax



### resource\_parameters::=



**password\_parameters::=****Purpose**

To create a profile. A **profile** is a set of limits on database resources. If you assign the profile to a user, that user cannot exceed these limits.

**Prerequisites**

You must have `CREATE PROFILE` system privilege.

To specify resource limits for a user, you must:

- Enable resource limits dynamically with the `ALTER SYSTEM` statement (see "[ALTER SYSTEM](#)" on page 7-102) or with the initialization parameter `RESOURCE_LIMIT`. (This parameter does not apply to password resources. Password resources are always enabled.)
- Create a profile that defines the limits using the `CREATE PROFILE` statement.
- Assign the profile to the user using the `CREATE USER` or `ALTER USER` statement (see "[CREATE USER](#)" on page 7-451 and "[ALTER USER](#)" on page 7-193).

**Keywords and Parameters**


---

*profile* is the name of the profile to be created. Use profiles to limit the database resources available to a user for a single call or a single session.

---

Oracle enforces resource limits in the following ways:

- If a user exceeds the `CONNECT_TIME` or `IDLE_TIME` session resource limit, Oracle rolls back the current transaction and ends the session. When the user process next issues a call, Oracle returns an error.
- If a user attempts to perform an operation that exceeds the limit for other session resources, Oracle aborts the operation, rolls back the current statement, and immediately returns an error. The user can then commit or roll back the current transaction, and must then end the session.
- If a user attempts to perform an operation that exceeds the limit for a single call, Oracle aborts the operation, rolls back the current statement, and returns an error, leaving the current transaction intact.

---

**Notes:**

- You can use fractions of days for all parameters that limit time, with days as units. For example, 1 hour is 1/24 and 1 minute is 1/1440.
- You can specify resource limits for users regardless of whether the resource limits are enabled. However, Oracle does not enforce the limits until you enable them.

UNLIMITED

When specified with a resource parameter, indicates that a user assigned this profile can use an unlimited amount of this resource. When specified with a password parameter, indicates that no limit has been set for the parameter.

DEFAULT

omits a limit for this resource in this profile. A user assigned this profile is subject to the limit for this resource specified in the `DEFAULT` profile. The `DEFAULT` profile initially defines unlimited resources. You can change those limits with the `ALTER PROFILE` statement.

Any user who is not explicitly assigned a profile is subject to the limits defined in the `DEFAULT` profile. Also, if the profile that is explicitly assigned to a user omits limits for some resources or specifies `DEFAULT` for some limits, the user is subject to the limits on those resources defined by the `DEFAULT` profile.

*resource\_parameters*

`SESSIONS_PER_USER`

limits a user to *integer* concurrent sessions.

`CPU_PER_SESSION`

limits the CPU time for a session, expressed in hundredth of seconds.

`CPU_PER_CALL`

limits the CPU time for a call (a parse, execute, or fetch), expressed in hundredths of seconds.

`CONNECT_TIME`

limits the total elapsed time of a session, expressed in minutes.

`IDLE_TIME`

limits periods of continuous inactive time during a session, expressed in minutes. Long-running queries and other operations are not subject to this limit.

## CREATE PROFILE

---

LOGICAL_READS_PER_SESSION	specifies the number of data blocks read in a session, including blocks read from memory and disk.
LOGICAL_READS_PER_CALL	specifies the number of data blocks read for a call to process a SQL statement (a parse, execute, or fetch).
PRIVATE_SGA	specifies the amount of private space a session can allocate in the shared pool of the system global area (SGA), expressed in bytes. Use K or M to specify this limit in kilobytes or megabytes.  <b>Note:</b> This limit applies only if you are using multi-threaded server architecture. The private space for a session in the SGA includes private SQL and PL/SQL areas, but not shared SQL and PL/SQL areas.
COMPOSITE_LIMIT	specifies the total resources cost for a session, expressed in <i>service units</i> . Oracle calculates the total service units as a weighted sum of CPU_PER_SESSION, CONNECT_TIME, LOGICAL_READS_PER_SESSION, and PRIVATE_SGA.  For information on how to specify the weight for each session resource, see " <a href="#">ALTER RESOURCE COST</a> " on page 7-73.
<i>password_parameters</i>	For a detailed description and explanation of how to use password management and protection, see <i>Oracle8i Administrator's Guide</i> .
FAILED_LOGIN_ATTEMPTS	specifies the number of failed attempts to log in to the user account before the account is locked.
PASSWORD_LIFE_TIME	limits the number of days the same password can be used for authentication. The password expires if it is not changed within this period, and further connections are rejected.
PASSWORD_REUSE_TIME	specifies the number of days before which a password cannot be reused. If you set PASSWORD_REUSE_TIME to an integer value, then you must set PASSWORD_REUSE_MAX to UNLIMITED.
PASSWORD_REUSE_MAX	specifies the number of password changes required before the current password can be reused. If you set PASSWORD_REUSE_MAX to an integer value, then you must set PASSWORD_REUSE_TIME to UNLIMITED.
PASSWORD_LOCK_TIME	specifies the number of days an account will be locked after the specified number of consecutive failed login attempts.
PASSWORD_GRACE_TIME	specifies the number of days after the grace period begins during which a warning is issued and login is allowed. If the password is not changed during the grace period, the password expires.
PASSWORD_VERIFY_FUNCTION	allows a PL/SQL password complexity verification script to be passed as an argument to the CREATE PROFILE statement. Oracle provides a default script, but you can create your own routine or use third-party software instead.  <i>function</i> is the name of the password complexity verification routine.

NULL

indicates that no password verification is performed.

**Restrictions on password parameters:**

- If `PASSWORD_REUSE_TIME` is set to an integer value, `PASSWORD_REUSE_MAX` must be set to `UNLIMITED`. If `PASSWORD_REUSE_MAX` is set to an integer value, `PASSWORD_REUSE_TIME` must be set to `UNLIMITED`.
- If both `PASSWORD_REUSE_TIME` and `PASSWORD_REUSE_MAX` are set to `UNLIMITED`, then Oracle uses neither of these password resources.
- If `PASSWORD_REUSE_MAX` is set to `DEFAULT` and `PASSWORD_REUSE_TIME` is set to `UNLIMITED`, then Oracle uses the `PASSWORD_REUSE_MAX` value defined in the `DEFAULT` profile.
- If `PASSWORD_REUSE_TIME` is set to `DEFAULT` and `PASSWORD_REUSE_MAX` is set to `UNLIMITED`, then Oracle uses the `PASSWORD_REUSE_TIME` value defined in the `DEFAULT` profile.
- If both `PASSWORD_REUSE_TIME` and `PASSWORD_REUSE_MAX` are set to `DEFAULT`, then Oracle uses whichever value is defined in the `DEFAULT` profile.

**Examples**

The following statement creates the profile `SYSTEM_MANAGER`:

```
CREATE PROFILE system_manager
  LIMIT SESSIONS_PER_USER      UNLIMITED
  CPU_PER_SESSION              UNLIMITED
  CPU_PER_CALL                  3000
  CONNECT_TIME                  45
  LOGICAL_READS_PER_SESSION    DEFAULT
  LOGICAL_READS_PER_CALL        1000
  PRIVATE_SGA                   15K
  COMPOSITE_LIMIT               5000000;
```

If you then assign the `SYSTEM_MANAGER` profile to a user, the user is subject to the following limits in subsequent sessions:

- The user can have any number of concurrent sessions.
- In a single session, the user can consume an unlimited amount of CPU time.
- A single call made by the user cannot consume more than 30 seconds of CPU time.
- A single session cannot last for more than 45 minutes.
- In a single session, the number of data blocks read from memory and disk is subject to the limit specified in the `DEFAULT` profile.

- A single call made by the user cannot read more than 1000 data blocks from memory and disk.
- A single session cannot allocate more than 15 kilobytes of memory in the SGA.
- In a single session, the total resource cost cannot exceed 5 million service units. The formula for calculating the total resource cost is specified by the ALTER RESOURCE COST statement.
- Since the SYSTEM\_MANAGER profile omits a limit for IDLE\_TIME and for password limits, the user is subject to the limits on these resources specified in the DEFAULT profile.

The following statement creates the profile PROF:

```
CREATE PROFILE prof
  LIMIT PASSWORD_REUSE_MAX DEFAULT
        PASSWORD_REUSE_TIME UNLIMITED;
```


The following statement creates profile MYPROFILE with password profile limits values set:

```
CREATE PROFILE myprofile LIMIT
  FAILED_LOGIN_ATTEMPTS 5
  PASSWORD_LIFE_TIME 60
  PASSWORD_REUSE_TIME 60
  PASSWORD_REUSE_MAX UNLIMITED
  PASSWORD_VERIFY_FUNCTION verify_function
  PASSWORD_LOCK_TIME 1/24
  PASSWORD_GRACE_TIME 10;
```



## CREATE ROLE

### Syntax



### Purpose

To create a **role**, which is a set of privileges that can be granted to users or to other roles. You can use roles to administer database privileges. You can add privileges to a role and then grant the role to a user. The user can then enable the role and exercise the privileges granted by the role.

A role contains all privileges granted to the role and all privileges of other roles granted to it. A new role is initially empty. You add privileges to a role with the GRANT statement. For information on granting roles, see "[GRANT system\\_privileges\\_and\\_roles](#)" on page 7-519. For information on enabling roles, see "[ALTER USER](#)" on page 7-193.

When you create a role that is NOT IDENTIFIED or is IDENTIFIED EXTERNALLY or BY *password*, Oracle grants you the role with ADMIN OPTION. However, when you create a role IDENTIFIED GLOBALLY, Oracle does not grant you the role.

For information on modifying a role, see "[ALTER ROLE](#)" on page 7-76. For information on removing a role from the database, see "[DROP ROLE](#)" on page 7-495. For information on enabling and disabling roles for the current session, see "[SET ROLE](#)" on page 7-600. For a detailed description and explanation of using global roles, see *Oracle8i Distributed Database Systems*.

### Prerequisites

You must have CREATE ROLE system privilege.

## Keywords and Parameters

<i>role</i>	is the name of the role to be created. Oracle recommends that the role contain at least one single-byte character regardless of whether the database character set also contains multibyte characters.  Some roles are defined by SQL scripts provided on your distribution media. For a list of these predefined roles, see " <a href="#">GRANT system_privileges_and_roles</a> " on page 7-519.
NOT IDENTIFIED	indicates that this role is authorized by the database and that no password is required to enable the role.
IDENTIFIED	indicates that a user must be authorized by the specified method before the role is enabled with the SET ROLE statement:
BY <i>password</i>	creates a <b>local user</b> and indicates that the user must specify the password to Oracle when enabling the role. The password can contain only single-byte characters from your database character set regardless of whether this character set also contains multibyte characters.
EXTERNALLY	creates an <b>external user</b> and indicates that a user must be authorized by an external service (such as an operating system or third-party service) before enabling the role.  Depending on the operating system, the user may have to specify a password to the operating system before the role is enabled.
GLOBALLY	creates a <b>global user</b> and indicates that a user must be authorized to use the role by the enterprise directory service before the role is enabled with the SET ROLE statement, or at login.

If you omit both the NOT IDENTIFIED clause and the IDENTIFIED clause, the role defaults to NOT IDENTIFIED.

---

## Examples

The following statement creates global role VENDOR:

```
CREATE ROLE vendor IDENTIFIED GLOBALLY;
```


The following statement creates the role TELLER:

```
CREATE ROLE teller  
    IDENTIFIED BY cashflow;
```

Users who are subsequently granted the TELLER role must specify the password CASHFLOW to enable the role with the SET ROLE statement.

## CREATE ROLLBACK SEGMENT

### Syntax



**storage\_clause:** See ["storage\\_clause"](#) on page 7-605.

### Purpose

To create a rollback segment. A **rollback segment** is an object that Oracle uses to store data necessary to reverse, or undo, changes made by transactions.

For information on altering a rollback segment, see ["ALTER ROLLBACK SEGMENT"](#) on page 7-78. For information on removing a rollback segment, see ["DROP ROLLBACK SEGMENT"](#) on page 7-496.

### Prerequisites

You must have CREATE ROLLBACK SEGMENT system privilege.

### Keyword and Parameters

<b>PUBLIC</b>	specifies that the rollback segment is public and is available to any instance. If you omit this clause, the rollback segment is private and is available only to the instance naming it in its initialization parameter <code>ROLLBACK_SEGMENTS</code> .
<i>rollback_segment</i>	is the name of the rollback segment to be created.
<b>TABLESPACE</b>	identifies the tablespace in which the rollback segment is created. If you omit this clause, Oracle creates the rollback segment in the SYSTEM tablespace.  <b>Restriction:</b> You cannot create a rollback segment in a tablespace that is system managed (that is, during creation you specified <code>EXTENT MANAGEMENT LOCAL AUTOALLOCATE</code> ). See <a href="#">"CREATE TABLESPACE"</a> on page 4-419.

**Notes:**

- A tablespace can have multiple rollback segments. Generally, multiple rollback segments improve performance.
- The tablespace must be online for you to add a rollback segment to it.
- When you create a rollback segment, it is initially offline. To make it available for transactions by your Oracle instance, bring it online using the `ALTER ROLLBACK SEGMENT` statement. To bring it online automatically whenever you start up the database, add the segment's name to the value of the `ROLLBACK_SEGMENTS` initialization parameter.

---

**See also:** *Oracle8i Administrator's Guide* for more information on creating rollback segments and making them available.

*storage\_clause* specifies the characteristics for the rollback segment. See the "[storage\\_clause](#)" on page 7-605.

**Notes:**

- The `OPTIMAL` parameter of the *storage\_clause* is of particular interest, because it applies only to rollback segments.
- You cannot specify the `PCTINCREASE` parameter of the *storage\_clause* with `CREATE ROLLBACK SEGMENT`.

---

## Examples

The following statement creates a rollback segment with default storage values in the system tablespace:


```
CREATE ROLLBACK SEGMENT rbs_2
  TABLESPACE system;
```

The above statement is equivalent to the following:

```
CREATE ROLLBACK SEGMENT rbs_2
  TABLESPACE system
  STORAGE
  ( INITIAL 10K
    NEXT 10K
    MAXEXTENTS UNLIMITED);
```

## CREATE SCHEMA

### Syntax



### Purpose

To create multiple tables and views and perform multiple grants in a single transaction.

To execute a `CREATE SCHEMA` statement, Oracle executes each included statement. If all statements execute successfully, Oracle commits the transaction. If any statement results in an error, Oracle rolls back all the statements.

---



---

**Note:** This statement does not actually create a schema. Oracle automatically creates a schema when you create a user (see ["CREATE USER"](#) on page 7-451). This statement lets you populate your schema with tables and views and grant privileges on those objects without having to issue multiple SQL statements in multiple transactions.

---



---

### Prerequisites

The `CREATE SCHEMA` statement can include `CREATE TABLE`, `CREATE VIEW`, and `GRANT` statements. To issue a `CREATE SCHEMA` statement, you must have the privileges necessary to issue the included statements.

### Keyword and Parameters

<i>schema</i>	is the name of the schema. The schema name must be the same as your Oracle username.
<i>create_table_statement</i>	is a <code>CREATE TABLE</code> statement to be issued as part of this <code>CREATE SCHEMA</code> statement. See <a href="#">"CREATE TABLE"</a> on page 4-381. Do not end this statement with a semicolon (or other terminator character).

---

<i>create_view_statement</i>	is a <code>CREATE VIEW</code> statement to be issued as part of this <code>CREATE SCHEMA</code> statement. See " <a href="#">CREATE VIEW</a> " on page 7-456. Do not end this statement with a semicolon (or other terminator character).
<i>grant_statement</i>	is a <code>GRANT object_privileges</code> statement to be issued as part of this <code>CREATE SCHEMA</code> statement. See " <a href="#">GRANT object_privileges</a> " on page 7-532. Do not end this statement with a semicolon (or other terminator character).

The `CREATE SCHEMA` statement supports the syntax of these statements only as defined by standard SQL, rather than the complete syntax supported by Oracle.

The order in which you list the `CREATE TABLE`, `CREATE VIEW`, and `GRANT` statements is unimportant. The statements within a `CREATE SCHEMA` statement can reference existing objects or objects you create in other statements within the same `CREATE SCHEMA` statement.

**Restriction:** The syntax of the *parallel\_clause* is allowed for a `CREATE TABLE` statement in `CREATE SCHEMA`, but parallelism is **not** used when creating the objects.

**See also:** The *parallel\_clause* of "[CREATE TABLE](#)" on page 4-381.

---


### Example

The following statement creates a schema named `BLAIR` for the user `BLAIR`, creates the table `SOX`, creates the view `RED_SOX`, and grants `SELECT` privilege on the `RED_SOX` view to the user `WAITES`.

```
CREATE SCHEMA AUTHORIZATION blair
  CREATE TABLE sox
    (color VARCHAR2(10) PRIMARY KEY, quantity NUMBER)
  CREATE VIEW red_sox
    AS SELECT color, quantity FROM sox WHERE color = 'RED'
  GRANT select ON red_sox TO waites;
```

## CREATE SEQUENCE

### Syntax



### Purpose

To create a sequence. A **sequence** is a database object from which multiple users may generate unique integers. You can use sequences to automatically generate primary key values.

When a sequence number is generated, the sequence is incremented, independent of the transaction committing or rolling back. If two users concurrently increment the same sequence, the sequence numbers each user acquires may have gaps because sequence numbers are being generated by the other user. One user can never acquire the sequence number generated by another user. Once a sequence value is generated by one user, that user can continue to access that value regardless of whether the sequence is incremented by another user.

Sequence numbers are generated independently of tables, so the same sequence can be used for one or for multiple tables. It is possible that individual sequence numbers will appear to be skipped, because they were generated and used in a transaction that ultimately rolled back. Additionally, a single user may not realize that other users are drawing from the same sequence.

Once a sequence is created, you can access its values in SQL statements with the `CURRVAL` pseudocolumn (which returns the current value of the sequence) or the `NEXTVAL` pseudocolumn (which increments the sequence and returns the new value).

**See Also:**

- ["Pseudocolumns"](#) on page 2-56 for more information on using the above pseudocolumns.
- ["ALTER SEQUENCE"](#) on page 7-81 or ["DROP SEQUENCE"](#) on page 7-497 for information on modifying or dropping a sequence.

### Prerequisites

To create a sequence in your own schema, you must have `CREATE SEQUENCE` privilege.

To create a sequence in another user's schema, you must have `CREATE ANY SEQUENCE` privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema to contain the sequence. If you omit <i>schema</i> , Oracle creates the sequence in your own schema.
<i>sequence</i>	is the name of the sequence to be created.



If you specify none of the following clauses, you create an ascending sequence that starts with 1 and increases by 1 with no upper limit. Specifying only `INCREMENT BY -1` creates a descending sequence that starts with -1 and decreases with no lower limit.

- **To create a sequence that increments without bound**, for ascending sequences, omit the `MAXVALUE` parameter or specify `NOMAXVALUE`. For descending sequences, omit the `MINVALUE` parameter or specify the `NOMINVALUE`.
- **To create a sequence that stops at a predefined limit**, for an ascending sequence, specify a value for the `MAXVALUE` parameter. For a descending sequence, specify a value for the `MINVALUE` parameter. Also specify the `NOCYCLE`. Any attempt to generate a sequence number once the sequence has reached its limit results in an error.
- **To create a sequence that restarts after reaching a predefined limit**, specify values for both the `MAXVALUE` and `MINVALUE` parameters. Also specify the `CYCLE`. If you do not specify `MINVALUE`, then it defaults to `NOMINVALUE` (that is, the value 1).

`INCREMENT BY` specifies the interval between sequence numbers. This integer value can be any positive or negative integer, but it cannot be 0. This value can have 28 or fewer digits. The absolute of this value must be less than the difference of `MAXVALUE` and `MINVALUE`. If this value is negative, then the sequence descends. If the increment is positive, then the sequence ascends. If you omit this clause, the interval defaults to 1.

`START WITH` specifies the first sequence number to be generated. Use this clause to start an ascending sequence at a value greater than its minimum or to start a descending sequence at a value less than its maximum. For ascending sequences, the default value is the sequence's minimum value. For descending sequences, the default value is the sequence's maximum value. This integer value can have 28 or fewer digits.

---

**Note:** This value is not necessarily the value to which an ascending cycling sequence cycles after reaching its maximum or minimum value.

---

`MAXVALUE` specifies the maximum value the sequence can generate. This integer value can have 28 or fewer digits. `MAXVALUE` must be equal to or greater than `START WITH` and must be greater than `MINVALUE`.

`NOMAXVALUE` specifies a maximum value of  $10^{27}$  for an ascending sequence or -1 for a descending sequence. This is the default.

`MINVALUE` specifies the sequence's minimum value. This integer value can have 28 or fewer digits. `MINVALUE` must be less than or equal to `START WITH` and must be less than `MAXVALUE`.

`NOMINVALUE` specifies a minimum value of 1 for an ascending sequence or  $-(10^{26})$  for a descending sequence. This is the default.

`CYCLE` specifies that the sequence continues to generate values after reaching either its maximum or minimum value. After an ascending sequence reaches its maximum value, it generates its minimum value. After a descending sequence reaches its minimum, it generates its maximum.

`NOCYCLE` specifies that the sequence cannot generate more values after reaching its maximum or minimum value. This is the default.

<b>CACHE</b>	<p>specifies how many values of the sequence Oracle preallocates and keeps in memory for faster access. This integer value can have 28 or fewer digits. The minimum value for this parameter is 2. For sequences that cycle, this value must be less than the number of values in the cycle. You cannot cache more values than will fit in a given cycle of sequence numbers. Therefore, the maximum value allowed for <b>CACHE</b> must be less than the value determined by the following formula:</p> $(\text{CEIL}(\text{MAXVALUE} - \text{MINVALUE})) / \text{ABS}(\text{INCREMENT})$ <p>If a system failure occurs, all cached sequence values that have not been used in committed DML statements are lost. The potential number of lost values is equal to the value of the <b>CACHE</b> parameter.</p>
<b>NOCACHE</b>	<p>specifies that values of the sequence are not preallocated.</p>
<p>If you omit both <b>CACHE</b> and <b>NOCACHE</b>, Oracle caches 20 sequence numbers by default.</p>	
<b>ORDER</b>	<p>guarantees that sequence numbers are generated in order of request. You may want to use this clause if you are using the sequence numbers as timestamps. Guaranteeing order is usually not important for sequences used to generate primary keys.</p> <p><b>ORDER</b> is necessary only to guarantee ordered generation if you are using Oracle with the Parallel Server option in parallel mode. If you are using exclusive mode, sequence numbers are always generated in order.</p>
<b>NOORDER</b>	<p>does not guarantee sequence numbers are generated in order of request. This is the default.</p>

---

### Example

The following statement creates the sequence **ESEQ**:

```
CREATE SEQUENCE eseq
  INCREMENT BY 10;
```

The first reference to **ESEQ.NEXTVAL** returns 1. The second returns 11. Each subsequent reference will return a value 10 greater than the one previous.

## CREATE SNAPSHOT

In Oracle8i, "snapshots" are synonymous with "materialized views." Please see ["CREATE MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-318.


---

## CREATE SNAPSHOT LOG

In Oracle8i, "snapshots" are synonymous with "materialized views." Please see ["CREATE MATERIALIZED VIEW LOG / SNAPSHOT LOG"](#) on page 7-333.

## CREATE SYNONYM

### Syntax



### Purpose

To create a synonym. A *synonym* is an alternative name for a table, view, sequence, procedure, stored function, package, materialized view, Java class schema object, or another synonym. For general information on synonyms, see *Oracle8i Concepts*.

Synonyms provide both data independence and location transparency. Synonyms permit applications to function without modification regardless of which user owns the table or view and regardless of which database holds the table or view.

[Table 4-4](#) lists the SQL statements in which you can refer to synonyms.

**Table 7-4 Using Synonyms**

DML Statements	DDL Statements
SELECT	AUDIT
INSERT	NOAUDIT
UPDATE	GRANT
DELETE	REVOKE
EXPLAIN PLAN	COMMENT
LOCK TABLE	

### Prerequisites

To create a private synonym in your own schema, you must have `CREATE SYNONYM` system privilege.

To create a private synonym in another user's schema, you must have `CREATE ANY SYNONYM` system privilege.

To create a PUBLIC synonym, you must have CREATE PUBLIC SYNONYM system privilege.

## Keywords and Parameters

---

PUBLIC	<p>creates a public synonym. Public synonyms are accessible to all users.</p> <p>Oracle uses a public synonym only when resolving references to an object if the object is not prefaced by a schema and the object is not followed by a database link.</p> <p>If you omit this clause, the synonym is private and is accessible only within its schema. A private synonym name must be unique in its schema.</p>
<i>schema</i>	<p>is the schema to contain the synonym. If you omit <i>schema</i>, Oracle creates the synonym in your own schema. You cannot specify a schema for the synonym if you have specified PUBLIC.</p>
<i>synonym</i>	<p>is the name of the synonym to be created.</p> <hr/> <p><b>CAUTION:</b> The functional maximum length of the synonym name is 32 bytes. Names longer than 30 bytes are permitted for Java functionality only. If you specify a name longer than 30 bytes, Oracle encrypts the name and places a representation of the encryption in the data dictionary. The actual encryption is not accessible, and you cannot use either your original specification or the data dictionary representation as the synonym name.</p> <hr/>
FOR	<p>identifies the object for which the synonym is created. If you do not qualify object with <i>schema</i>, Oracle assumes that the schema object is in your own schema. The schema object can be of the following types:</p> <ul style="list-style-type: none"><li>■ table or object table</li><li>■ view or object view</li><li>■ sequence</li><li>■ stored procedure, function, or package</li><li>■ materialized view</li><li>■ Java class schema object</li><li>■ synonym</li></ul> <p>The schema object need not currently exist and you need not have privileges to access the object.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ The schema object cannot be contained in a package.</li><li>■ You cannot create a synonym for an object type.</li></ul>

---

***dblink***

You can use a complete or partial *dblink* to create a synonym for a schema object on a remote database where the object is located. If you specify *dblink* and omit *schema*, the synonym refers to an object in the schema specified by the database link. Oracle Corporation recommends that you specify the schema containing the object in the remote database.

**See Also:**

- ["Referring to Objects in Remote Databases"](#) on page 2-79 for more information on referring to database links.
- ["CREATE DATABASE LINK"](#) on page 7-273 for more information on creating database links.

If you omit *dblink*, Oracle assumes the object is located on the local database.

**Restriction:** You cannot specify *dblink* for a Java class synonym.

---

## Examples

Oracle attempts to resolve references to objects at the schema level before resolving them at the PUBLIC synonym level. For example, assume the schemas SCOTT and BLAKE each contain tables named DEPT and the user SYSTEM creates a PUBLIC synonym named DEPT for BLAKE . DEPT. If the user SCOTT then issues the following statement, Oracle returns rows from SCOTT . DEPT:

```
SELECT * FROM dept;
```

To retrieve rows from BLAKE . DEPT, the user SCOTT must preface DEPT with the schema name:

```
SELECT * FROM blake.dept;
```

If the user ADAM's schema does not contain an object named DEPT, then ADAM can access the DEPT table in BLAKE's schema by using the public synonym DEPT:

```
SELECT * FROM dept;
```

To define the synonym MARKET for the table MARKET\_RESEARCH in the schema SCOTT, issue the following statement:

```
CREATE SYNONYM market
FOR scott.market_research;
```

To create a PUBLIC synonym for the EMP table in the schema SCOTT on the remote SALES database, you could issue the following statement:

```
CREATE PUBLIC SYNONYM emp
FOR scott.emp@sales;
```


A synonym may have the same name as the base table, provided the base table is contained in another schema.



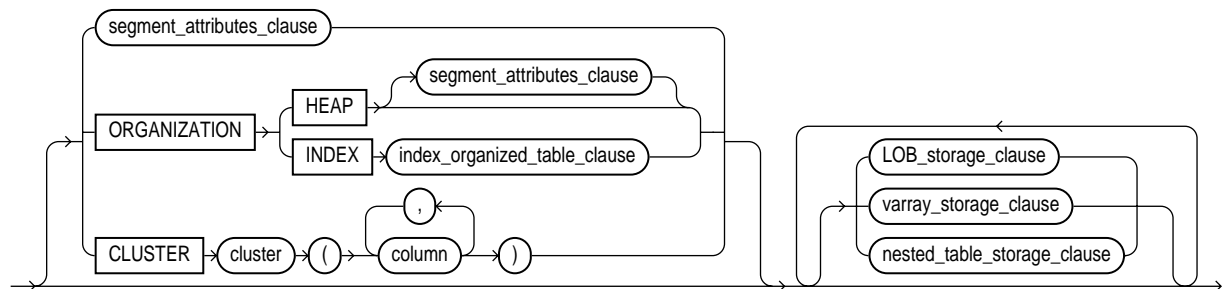


## CREATE TABLE

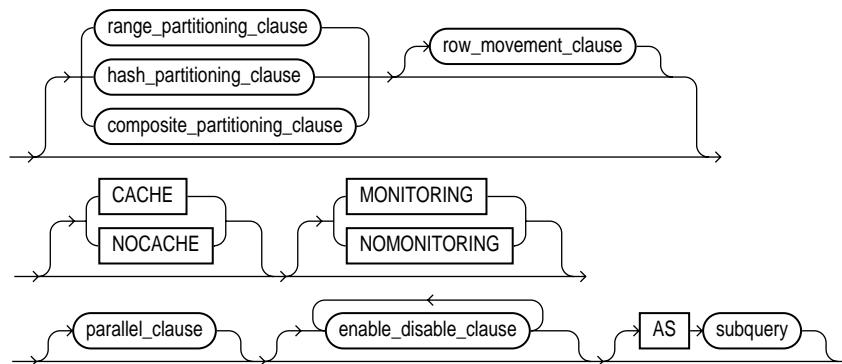
### object\_properties::=



### physical\_properties::=







### table\_properties::=




**subquery::=** See ["SELECT and Subqueries"](#) on page 7-569.

**table\_constraint, column\_constraint, table\_ref\_constraint, column\_ref\_constraint, constraint\_state:** See the ["constraint\\_clause"](#) on page 7-233


**OID\_clause::=****OID\_index\_clause::=****segment\_attributes\_clause::=****row\_movement\_clause::=**

**physical\_attributes\_clause::=**




**storage\_clause:** See the "[storage\\_clause](#)" on page 7-605.

**index\_organized\_table\_clause::=**


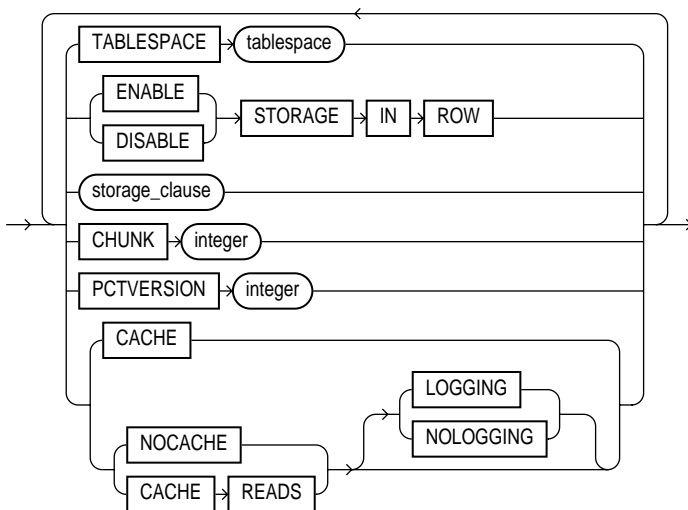



**compression\_clause::=**



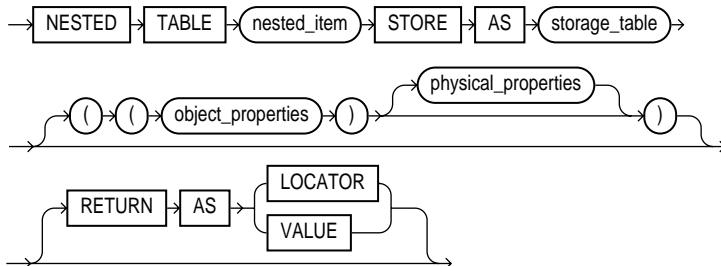
**index\_organized\_overflow\_clause::=**



**LOB\_storage\_clause::=****LOB\_parameters::=****varray\_storage\_clause::=**

## CREATE TABLE


### nested\_table\_storage\_clause::=



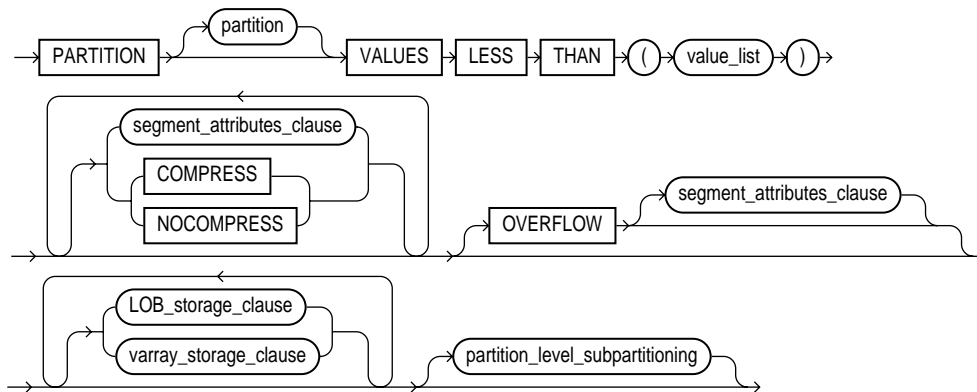
### range\_partitioning\_clause::=



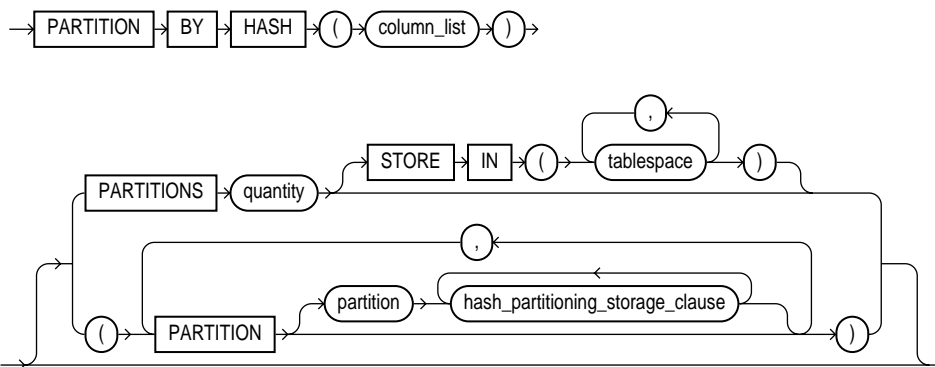


### composite\_partitioning\_clause::=




### partition\_definition::=




**subpartition\_clause::=****partition\_level\_subpartitioning::=****hash\_partitioning\_clause::=**

## CREATE TABLE

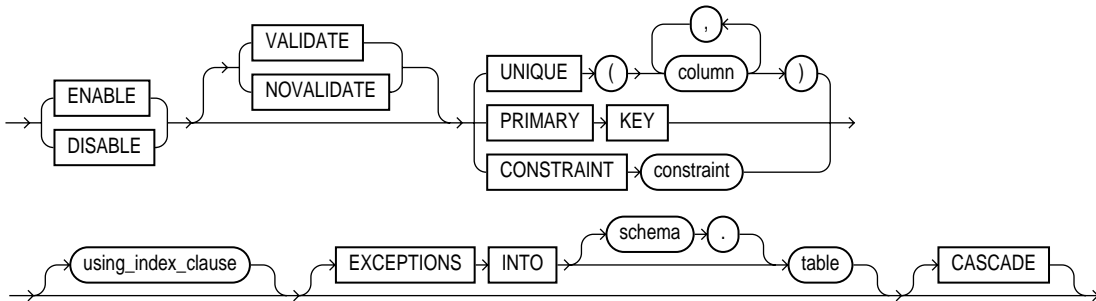
### hash\_partitioning\_storage\_clause::=



### parallel\_clause::=




### enable\_disable\_clause::=






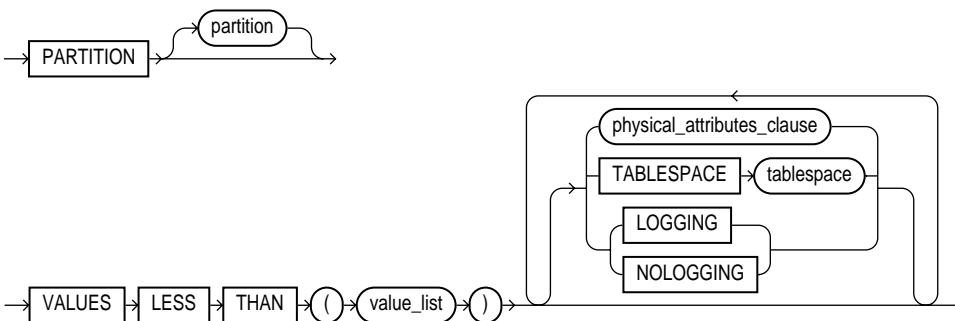
**using\_index\_clause::=**



**global\_index\_clause::=**



**global\_partition\_clause::=**



## Purpose

To create a **relational table**, the basic structure to hold user data.

To create an **object table** or a table that uses an object type for a column definition. An object table is a table explicitly defined to hold object instances of a particular type.

You can also create an object type and then use it in a column when creating a relational table.

Tables are created with no data unless a query is specified. You can add rows to a table with the `INSERT` statement. After creating a table, you can define additional columns, partitions, and integrity constraints with the `ADD` clause of the `ALTER TABLE` statement. You can change the definition of an existing column or partition with the `MODIFY` clause of the `ALTER TABLE` statement.

**See also:** *Oracle8i Application Developer's Guide - Fundamentals* and "[CREATE TYPE](#)" on page 7-437 for more information about creating objects.

## Prerequisites

To create a relational table in your own schema, you must have system privilege. To create a table in another user's schema, you must have `CREATE ANY TABLE` system privilege. Also, the owner of the schema to contain the table must have either space quota on the tablespace to contain the table or `UNLIMITED TABLESPACE` system privilege.

In addition to the table privileges above, to create a table that uses types, the owner of the table must have the `EXECUTE` object privilege in order to access all types referenced by the table, or you must have the `EXECUTE ANY TYPE` system privilege. These privileges must be granted explicitly and not acquired through a role.

Additionally, if the table owner intends to grant access to the table to other users, the owner must have been granted the `EXECUTE` privileges to the referenced types with the `GRANT OPTION`, or have the `EXECUTE ANY TYPE` system privilege with the `ADMIN OPTION`. Without these privileges, the table owner has insufficient privileges to grant access on the table to other users.

To enable a `UNIQUE` or `PRIMARY KEY` constraint, you must have the privileges necessary to create an index on the table. You need these privileges because Oracle creates an index on the columns of the unique or primary key in the schema containing the table.

**See Also:**

- ["CREATE INDEX"](#) on page 7-291.
- *Oracle8i Application Developer's Guide - Fundamentals* for more information about the privileges required to create tables using types.

**Keywords and Parameters**


---

GLOBAL TEMPORARY	<p>specifies that the table is temporary and that its definition is visible to all sessions. The data in a temporary table is visible only to the session that inserts the data into the table.</p> <p>A temporary table has a definition that persists the same as the definitions of regular tables, but it contains either <b>session-specific</b> or <b>transaction-specific</b> data. You specify whether the data is session- or transaction-specific with the <code>ON COMMIT</code> keywords (below).</p> <p><b>See Also:</b> <i>Oracle8i Concepts</i> for more information on temporary tables.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ Temporary tables cannot be partitioned, index-organized, or clustered.</li> <li>■ You cannot specify any referential integrity (foreign key) constraints on temporary tables.</li> <li>■ Temporary tables cannot contain columns of nested table or varray type.</li> <li>■ You cannot specify the following clauses of the <i>LOB_storage_clause</i>: <code>TABLESPACE</code>, <i>storage_clause</i>, <code>LOGGING</code> or <code>NOLOGGING</code>, <code>MONITORING</code> or <code>NOMONITORING</code>, or <i>LOB_index_clause</i>.</li> <li>■ Parallel DML and parallel queries are not supported for temporary tables. (Parallel hints are ignored. Specification of the <i>parallel_clause</i> returns an error.)</li> <li>■ You cannot specify the <i>segment_attributes_clause</i>, <i>nested_table_storage_clause</i>, or <i>parallel_clause</i>.</li> <li>■ Distributed transactions are not supported for temporary tables.</li> </ul>
<i>schema</i>	is the schema to contain the table. If you omit <i>schema</i> , Oracle creates the table in your own schema.
<i>table</i>	is the name of the table (or object table) to be created.
OF <i>object_type</i>	<p>explicitly creates an object table of type <i>object_type</i>. The columns of an object table correspond to the top-level attributes of type <i>object_type</i>. Each row will contain an object instance, and each instance will be assigned a unique, system-generated object identifier (OID) when a row is inserted. If you omit <i>schema</i>, Oracle creates the object table in your own schema.</p> <p><b>See Also:</b> <a href="#">"CREATE TYPE"</a> on page 7-437 for more information about creating objects.</p>

---

---

	<p>Objects residing in an object table are referenceable.</p> <p><b>See Also:</b> "User-Defined Type Categories" on page 2-30, "User-Defined Functions" on page 4-118, "Expressions" on page 5-1, "CREATE TYPE" on page 7-437, and <i>Oracle8i Administrator's Guide</i> for more information about using REFS.</p>
<i>column</i>	<p>specifies the name of a column of the table.</p> <p>If you also specify <i>AS subquery</i>, you can omit <i>column</i> and <i>datatype</i> unless you are creating an index-organized table (IOT). If you specify <i>AS subquery</i> when creating an IOT, you must specify <i>column</i>, and you must omit <i>datatype</i>.</p> <p>The absolute maximum number of columns in a table is 1000. However, when you create an object table (or a relational table with columns of object, nested table, varray, or REF type), Oracle maps the columns of the user-defined types to relational columns, creating in effect "hidden columns" that count toward the 1000-column limit. For details on how Oracle calculates the total number of columns in such a table, please refer to <i>Oracle8i Administrator's Guide</i>.</p>
<i>attribute</i>	<p>specifies the qualified column name of an item in an object.</p>
<i>datatype</i>	<p>is the datatype of a column. Oracle-supplied datatypes are defined in "Datatypes" on page 2-9.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot specify a LOB column or a column of type VARRAY for a partitioned index-organized table. The datatypes for nonpartitioned index-organized tables are not restricted.</li><li>■ You can specify a column of type ROWID, but Oracle does not guarantee that the values in such columns are valid rowids.</li></ul> <p>You can omit <i>datatype</i>:</p> <ul style="list-style-type: none"><li>■ If you also specify <i>AS subquery</i>. (If you are creating an index-organized table and you specify <i>AS subquery</i>, you <b>must</b> omit the datatype.)</li><li>■ If the statement also designates the column as part of a foreign key in a referential integrity constraint. (Oracle automatically assigns to the column the datatype of the corresponding column of the referenced key of the referential integrity constraint.)</li></ul>
DEFAULT	<p>specifies a value to be assigned to the column if a subsequent INSERT statement omits a value for the column. The datatype of the expression must match the datatype of the column. The column must also be long enough to hold this expression. For the syntax of <i>expr</i>; see "Expressions" on page 5-1.</p> <p><b>Restriction:</b> A DEFAULT expression cannot contain references to other columns, the pseudocolumns CURRVAL, NEXTVAL, LEVEL, and ROWNUM, or date constants that are not fully specified.</p>

---

<i>table_ref_constraint</i>	These clauses let you further describe a column of type REF. The only difference between these clauses is that you specify <i>table_ref</i> from the table level, so you must identify the REF column or attribute you are defining. You specify <i>column_ref</i> after you have already identified the REF column or attribute.
<i>column_ref_constraint</i>	For syntax and description of these constraints, see the " <a href="#">constraint_clause</a> " on page 7-233.
<i>column_constraint</i>	defines an integrity constraint as part of the column definition. See the syntax description of <i>column_constraint</i> in the " <a href="#">constraint_clause</a> " on page 7-233.  You can create UNIQUE, PRIMARY KEY, and REFERENCES constraints on scalar attributes of object type columns. You can also create NOT NULL constraints on object type columns, and CHECK constraints that reference object type columns or any attribute of an object type column.
<i>table_constraint</i>	defines an integrity constraint as part of the table definition. See the syntax description of <i>table_constraint</i> in the " <a href="#">constraint_clause</a> " on page 7-233.  <b>Note:</b> You must specify a PRIMARY KEY constraint for an index-organized table, and it cannot be DEFERRABLE.
<hr/>	
<i>segment_attributes_clause:</i>	
<i>physical_attributes_clause</i>	specifies the value of the PCTFREE, PCTUSED, INITRANS, and MAXTRANS parameters and the storage characteristics of the table. <ul style="list-style-type: none"> <li>■ For a nonpartitioned table, each parameter and storage characteristic you specify determines the actual physical attribute of the segment associated with the table.</li> <li>■ For partitioned tables, the value you specify for the parameter or storage characteristic is the default physical attribute of the segments associated with all partitions specified in this CREATE statement (and in subsequent ALTER TABLE ... ADD PARTITION statements), unless you explicitly override that value in the PARTITION clause of the statement that creates the partition.</li> </ul>
PCTFREE	specifies the percentage of space in each data block of the table, object table OID index, or partition reserved for future updates to the table's rows. The value of PCTFREE must be a value from 0 to 99. A value of 0 allows the entire block to be filled by inserts of new rows. The default value is 10. This value reserves 10% of each block for updates to existing rows and allows inserts of new rows to fill a maximum of 90% of each block.  PCTFREE has the same function in the PARTITION description and in the statements that create and alter clusters, indexes, materialized views, and materialized view logs. The combination of PCTFREE and PCTUSED determines whether new rows will be inserted into existing data blocks or into new blocks.
PCTUSED	specifies the minimum percentage of used space that Oracle maintains for each data block of the table, object table OID index, or index-organized table overflow data segment. A block becomes a candidate for row insertion when its used space falls below PCTUSED. PCTUSED is specified as a positive integer from 0 to 99 and defaults to 40.

PCTUSED has the same function in the PARTITION description and in the statements that create and alter clusters, materialized views, and materialized view logs.

PCTUSED is not a valid table storage characteristic for an index-organized table (ORGANIZATION INDEX).

The sum of PCTFREE and PCTUSED must be equal to or less than 100. You can use PCTFREE and PCTUSED together to utilize space within a table more efficiently. For information on the performance effects of different values PCTUSED and PCTFREE, see *Oracle8i Designing and Tuning for Performance*.

**INITRANS**

specifies the initial number of transaction entries allocated within each data block allocated to the table, object table OID index, partition, LOB index segment, or overflow data segment. This value can range from 1 to 255 and defaults to 1. In general, you should not change the INITRANS value from its default.

Each transaction that updates a block requires a transaction entry in the block. The size of a transaction entry depends on your operating system.

This parameter ensures that a minimum number of concurrent transactions can update the block and helps avoid the overhead of dynamically allocating a transaction entry.

The INITRANS parameter serves the same purpose in the PARTITION description, clusters, indexes, materialized views, and materialized view logs as in tables. The minimum and default INITRANS value for a cluster or index is 2, rather than 1.

**MAXTRANS**

specifies the maximum number of concurrent transactions that can update a data block allocated to the table, object table OID index, partition, LOB index segment, or index-organized overflow data segment. This limit does not apply to queries. This value can range from 1 to 255 and the default is a function of the data block size. You should not change the MAXTRANS value from its default.

If the number of concurrent transactions updating a block exceeds the INITRANS value, Oracle dynamically allocates transaction entries in the block until either the MAXTRANS value is exceeded or the block has no more free space.

The MAXTRANS parameter serves the same purpose in the PARTITION description, clusters, materialized views, and materialized view logs as in tables.

**storage\_clause**

specifies the storage characteristics for the table, object table OID index, partition, LOB storage, LOB index segment, or index-organized table overflow data segment. This clause has performance ramifications for large tables. Storage should be allocated to minimize dynamic allocation of additional space. See the "[storage\\_clause](#)" on page 7-605.

**TABLESPACE**

specifies the tablespace in which Oracle creates the table, object table OID index, partition, LOB storage, LOB index segment, or index-organized table overflow data segment. If you omit TABLESPACE, then Oracle creates that item in the default tablespace of the owner of the schema containing the table.

---

For heap-organized tables with one or more LOB columns, if you omit the `TABLESPACE` clause for LOB storage, Oracle creates the LOB data and index segments in the tablespace where the table is created.

However, for an index-organized table with one or more LOB columns, if you omit `TABLESPACE`, the LOB data and index segments are created in the tablespace in which the primary key index segment of the index-organized table is created.

For nonpartitioned tables, the value specified for `TABLESPACE` is the actual physical attribute of the segment associated with the table. For partitioned tables, the value specified for `TABLESPACE` is the default physical attribute of the segments associated with all partitions specified in the `CREATE` statement (and on subsequent `ALTER TABLE . . . ADD PARTITION` statements), unless you specify `TABLESPACE` in the `PARTITION` description.

**See Also:** "[CREATE TABLESPACE](#)" on page 4-419 for more information on tablespaces.

LOGGING |  
NOLOGGING

specifies whether the creation of the table (and any indexes required because of constraints), partition, or LOB storage characteristics will be logged in the redo log file. The logging attribute of the table is independent of that of its indexes.

This attribute also specifies that subsequent Direct Loader (`SQL*Loader`) and direct-load `INSERT` operations against the table, partition, or LOB storage are logged (`LOGGING`) or not logged (`NOLOGGING`).

**For a table or table partition**, if you omit `LOGGING|NOLOGGING`, the logging attribute of the table or table partition defaults to the logging attribute of the tablespace in which it resides.

**For LOBs**, if you omit `LOGGING|NOLOGGING`,

- If you specify `CACHE`, then `LOGGING` is used (because you cannot have `CACHE NOLOGGING`).
- If you specify `NOCACHE` or `CACHE READS`, the logging attribute defaults to the logging attribute of the tablespace in which it resides.

`NOLOGGING` does not apply to LOBs that are stored inline with row data. That is, if you specify `NOLOGGING` for LOBs with values less than 4000 bytes and you have not disabled `STORAGE IN ROW`, Oracle ignores the `NOLOGGING` specification and treats the LOB data the same as other table data.

**For nonpartitioned tables**, the value specified for `LOGGING` is the actual physical attribute of the segment associated with the table. For partitioned tables, the logging attribute value specified is the default physical attribute of the segments associated with all partitions specified in the `CREATE` statement (and in subsequent `ALTER TABLE . . . ADD PARTITION` statements), unless you specify `LOGGING|NOLOGGING` in the `PARTITION` description.

In `NOLOGGING` mode, data is modified with minimal logging (to mark new extents `INVALID` and to record dictionary changes). When applied during media recovery, the extent invalidation records mark a range of blocks as logically corrupt, because the redo data is not fully logged. Therefore, if you cannot afford to lose this table, you should take a backup after the `NOLOGGING` operation.

The size of a redo log generated for an operation in `NOLOGGING` mode is significantly smaller than the log generated with the `LOGGING` attribute set.

If the database is run in `ARCHIVELOG` mode, media recovery from a backup taken before the `LOGGING` operation restores the table. However, media recovery from a backup taken before the `NOLOGGING` operation does not restore the table.

**See Also:** *Oracle8i Concepts* and *Oracle8i Administrator's Guide* for more information about logging and parallel DML.

RECOVERABLE |  
UNRECOVERABLE

These keywords are deprecated and have been replaced with `LOGGING` and `NOLOGGING`, respectively. Although `RECOVERABLE` and `UNRECOVERABLE` are supported for backward compatibility, Oracle Corporation strongly recommends that you use the `LOGGING` and `NOLOGGING` keywords.

**Restrictions:**

- You cannot specify `RECOVERABLE` for partitioned tables or LOB storage characteristics.
- You cannot specify `UNRECOVERABLE` for a partitioned or index-organized tables.
- You can specify `UNRECOVERABLE` only with `AS subquery`.

ORGANIZATION  
HEAP

specifies that the data rows of *table* are stored in no particular order. This is the default.

ORGANIZATION  
INDEX

specifies that *table* is created as an index-organized table. In an index-organized table, the data rows are held in an index defined on the primary key for the table.

*index\_organized\_  
table\_clause*

specifies that Oracle should maintain the table rows (both primary key column values and non-key column values) in a B\*-tree index built on the primary key. Index-organized tables are therefore best suited for primary key-based access and manipulation. An index-organized table is an alternative to

- A nonclustered table indexed on the primary key by using the `CREATE INDEX` statement
- A clustered table stored in an indexed cluster that has been created using the `CREATE CLUSTER` statement that maps the primary key for the table to the cluster key

**Restrictions:**

- You cannot specify a column of type `ROWID` for an index-organized table.
- A partitioned index-organized table cannot contain columns of LOB or varray type. (This restriction does not apply to nonpartitioned index-organized tables.)

---

**Note:** You must specify a primary key for an index-organized table, because the primary key uniquely identifies a row. The primary key cannot be `DEFERRABLE`. Use the primary key instead of the `rowid` for directly accessing index-organized rows.



<p>PCTTHRESHOLD <i>integer</i></p>	<p>specifies the percentage of space reserved in the index block for an index-organized table row. Any portion of the row that exceeds the specified threshold is stored in the overflow segment. PCTTHRESHOLD must be a value from 1 to 50.</p> <p><b>Restriction:</b></p> <ul style="list-style-type: none"> <li>■ PCTTHRESHOLD must be large enough to hold the primary key.</li> <li>■ You cannot specify PCTTHRESHOLD for individual partitions of an index-organized table.</li> </ul>
<p>OVERFLOW</p>	<p>specifies that index-organized table data rows exceeding the specified threshold are placed in the data segment listed in this clause.</p> <ul style="list-style-type: none"> <li>■ When you create an index-organized table, Oracle evaluates the maximum size of each column to estimate the largest possible row. If an overflow segment is needed but you have not specified OVERFLOW, Oracle raises an error and does not execute the CREATE TABLE statement. This checking function guarantees that subsequent DML operations on the index-organized table will not fail because an overflow segment is lacking.</li> <li>■ All physical attributes and storage characteristics you specify in this clause after the OVERFLOW keyword apply only to the overflow segment of the table. Physical attributes and storage characteristics for the index-organized table itself, default values for all its partitions, and values for individual partitions must be specified before this keyword.</li> <li>■ If the index-organized table contains one or more LOB columns, the LOBs will be stored out-of-line unless you specify OVERFLOW, even if they would otherwise be small enough to be stored inline.</li> </ul>
<p>INCLUDING <i>column_name</i></p>	<p>specifies a column at which to divide an index-organized table row into index and overflow portions. Oracle moves all primary key columns of an index-organized table to the beginning of the table (in their key order). <i>column_name</i> can be either the last primary-key column or any non-primary-key column. All non-primary-key columns that follow <i>column_name</i> are stored in the overflow data segment.</p> <p><b>Restriction:</b> You cannot specify this clause for individual partitions of an index-organized table.</p>
<p><i>compression_clause</i></p>	<p>enables or disables key compression.</p>
<p>COMPRESS</p>	<p>enables key compression, which eliminates repeated occurrence of primary key column values in index-organized tables. Use <i>integer</i> to specify the prefix length (number of prefix columns to compress).</p> <p>The valid range of prefix length values is from 1 to the number of primary key columns minus 1. The default prefix length is the number of primary key columns minus 1.</p> <p><b>Restriction:</b> At the partition level, you can specify COMPRESS, but you cannot specify the prefix length with <i>integer</i>.</p>
<p>NOCOMPRESS</p>	<p>disables key compression in index-organized tables. This is the default.</p>

---

<i>LOB_storage_clause</i>	<p>specifies the storage attributes of LOB data segments.</p> <ul style="list-style-type: none"> <li>■ For a nonpartitioned table (that is, when specified in the <i>physical_properties</i> clause without any of the partitioning clauses), this clause specifies the table's storage attributes of LOB data segments.</li> <li>■ For a partitioned table specified at the table level (that is, when specified in the <i>physical_properties</i> clause along with one of the partitioning clauses), this clause specifies the default storage attributes for LOB data segments associated with each partition or subpartition. These storage attributes apply to all partitions or subpartitions unless overridden by a <i>LOB_storage_clause</i> at the partition or subpartition level.</li> <li>■ For an individual partition of a partitioned table (that is, when specified as part of a <i>partition_definition</i>), this clause specifies the storage attributes of the data segments of that partition or the default storage attributes of any subpartitions of this partition. A partition-level <i>LOB_storage_clause</i> overrides a table-level <i>LOB_storage_clause</i>.</li> <li>■ For an individual subpartition of a partitioned table (that is, when specified as part of a <i>subpartition_clause</i>), this clause specifies the storage attributes of the data segments of this subpartition. A subpartition-level <i>LOB_storage_clause</i> overrides both partition-level and table-level <i>LOB_storage_clauses</i>.</li> </ul> <p><b>Restriction:</b> You cannot specify the <i>LOB_index_clause</i> if <i>table</i> is partitioned.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ <i>Oracle8i Application Developer's Guide - Large Objects (LOBs)</i> for detailed information about LOBs, including guidelines for creating gigabyte LOBs.</li> <li>■ "<a href="#">LOB Column Example</a>" on page 4-413</li> </ul>
<i>lob_item</i>	<p>is the LOB column name or LOB object attribute for which you are explicitly defining tablespace and storage characteristics that are different from those of the table. Oracle automatically creates a system-managed index for each <i>lob_item</i> you create.</p>
STORE AS	<p><i>lob_segname</i> specifies the name of the LOB data segment. You cannot use <i>lob_segname</i> if you specify more than one <i>lob_item</i>.</p>
<i>lob_parameters</i>	<p>ENABLE STORAGE IN ROW specifies that the LOB value is stored in the row (inline) if its length is less than approximately 4000 bytes minus system control information. This is the default.</p> <p><b>Restriction:</b> For an index-organized table, you cannot specify this parameter unless you have specified an OVERFLOW segment in the <i>index_organized_table_clause</i>.</p> <p>DISABLE STORAGE IN ROW specifies that the LOB value is stored outside of the row regardless of the length of the LOB value.</p>

The LOB locator is always stored in the row regardless of where the LOB value is stored. You cannot change the value of `STORAGE IN ROW` once it is set.

`CHUNK integer` specifies the number of bytes to be allocated for LOB manipulation. If *integer* is not a multiple of the database block size, Oracle rounds up (in bytes) to the next multiple. For example, if the database block size is 2048 and *integer* is 2050, Oracle allocates 4096 bytes (2 blocks). The maximum value is 32768 (32K), which is the largest Oracle block size allowed. The default `CHUNK` size is one Oracle database block.

You cannot change the value of `CHUNK` once it is set.

**Note:** The value of `CHUNK` must be less than or equal to the value of `NEXT` (either the default value or that specified in the *storage\_clause*). If `CHUNK` exceeds the value of `NEXT`, Oracle returns an error.

`PCTVERSION integer` is the maximum percentage of overall LOB storage space used for creating new versions of the LOB. The default value is 10, meaning that older versions of the LOB data are not overwritten until 10% of the overall LOB storage space is used.

`LOB_index_clause` This clause is deprecated as of Oracle8i. Oracle generates an index for each LOB column. Oracle names and manages the LOB indexes internally.

Although it is still possible for you to specify this clause, Oracle Corporation strongly recommends that you no longer do so.

For information on how Oracle manages LOB indexes in tables migrated from earlier versions, see *Oracle8i Migration*.

`varray_storage_clause` lets you specify separate storage characteristics for the LOB in which a varray will be stored. In addition, if you specify this clause, Oracle will always store the varray in a LOB, even if it is small enough to be stored inline.

- For a nonpartitioned table (that is, when specified in the *physical\_properties* clause without any of the partitioning clauses), this clause specifies the storage attributes of the varray's LOB data segments.
- For a partitioned table specified at the table level (that is, when specified in the *physical\_properties* clause along with one of the partitioning clauses), this clause specifies the default storage attributes for the varray's LOB data segments associated with each partition (or its subpartitions, if any).
- For an individual partition of a partitioned table (that is, when specified as part of a *partition\_definition*), this clause specifies the storage attributes of the varray's LOB data segments of that partition or the default storage attributes of the varray's LOB data segments of any subpartitions of this partition. A partition-level *varray\_storage\_clause* overrides a table-level *varray\_storage\_clause*.

- For an individual subpartition of a partitioned table (that is, when specified as part of a *subpartition\_clause*), this clause specifies the storage attributes of the varray's data segments of this subpartition. A subpartition-level *varray\_storage\_clause* overrides both partition-level and table-level *varray\_storage\_clauses*.

**Restriction:** You cannot specify the `TABLESPACE` parameter of *lob\_parameters* as part of this clause. The LOB tablespace for a varray defaults to the containing table's tablespace.

#### *nested\_table\_storage\_clause*

enables you to specify separate storage characteristics for a nested table, which in turn enables you to define the nested table as an index-organized table. The storage table is created in the same tablespace as its parent table (using the default storage characteristics) and stores the nested table values of the column for which it was created.

You must include this clause when creating a table with columns or column attributes whose type is a nested table. (Clauses within this clause that function the same way they function for parent object tables are not repeated here.)

#### **Restrictions:**

- You cannot specify this clause for a temporary table.
- You cannot specify the *OID\_clause*.
- You cannot specify `TABLESPACE` (as part of the *segment\_attributes\_clause*) for a nested table. The tablespace is always that of the parent table.
- At create time, you cannot specify (as part of *object\_properties*) a *table\_ref\_constraint*, *column\_ref\_constraint*, or referential constraint for the attributes of a nested table. However, you can modify a nested table to add such constraints using `ALTER TABLE`.
- You cannot query or perform DML statements on the storage table directly, but you can modify the nested table column storage characteristics by using the name of storage table in an `ALTER TABLE` statement. For information about modifying nested table column storage characteristics, see "[ALTER TABLE](#)" on page 7-123.

*nested\_item* is the name of a column (or a top-level attribute of the table's object type) whose type is a nested table.

*storage\_table* is the name of the table where the rows of *nested\_item* reside. For a nonpartitioned table, the storage table is created in the same schema and the same tablespace as the parent table. For a partitioned table, the storage table is created in the default tablespace of the schema.

**Restriction:** You cannot partition the storage table of a nested table.

You cannot query or perform DML statements on *storage\_table* directly, but you can modify its storage characteristics by specifying its name in an `ALTER TABLE` statement. For information about modifying nested table column storage characteristics, see "[ALTER TABLE](#)" on page 7-123.

---

RETURN AS	<p>specifies what Oracle returns as the result of a query.</p> <ul style="list-style-type: none"> <li>■ VALUE returns a copy of the nested table itself.</li> <li>■ LOCATOR returns a collection locator to the copy of the nested table.</li> </ul> <hr/> <p><b>Note:</b> The locator is scoped to the session and cannot be used across sessions. Unlike a LOB locator, the collection locator cannot be used to modify the collection instance.</p> <hr/>
CLUSTER	<p>If you do not specify the <i>segment_attributes_clause</i> or the <i>LOB_storage_clause</i>, the nested table is heap organized and is created with default storage characteristics.</p> <p>specifies that the table is to be part of <i>cluster</i>. The columns listed in this clause are the table columns that correspond to the cluster's columns. Generally, the cluster columns of a table are the column or columns that make up its primary key or a portion of its primary key.</p> <p><b>See Also:</b> "<a href="#">CREATE CLUSTER</a>" on page 7-254.</p> <p>Specify one column from the table for each column in the cluster key. The columns are matched by position, not by name.</p> <p>A clustered table uses the cluster's space allocation. Therefore, do not use the PCTFREE, PCTUSED, INITRANS, or MAXTRANS parameters, the TABLESPACE clause, or the <i>storage_clause</i> with the CLUSTER clause.</p> <p><b>Restriction:</b> Object tables cannot be part of a cluster.</p>
ON COMMIT	<p>can be specified only if you are creating a temporary table. This clause specifies whether the data in the temporary table persists for the duration of a transaction or a session.</p>
DELETE ROWS	<p>specifies that the temporary table is transaction specific (this is the default). Oracle will truncate the table (delete all its rows) after each commit.</p>
PRESERVE ROWS	<p>specifies that the temporary table is session specific. Oracle will truncate the table (delete all its rows) when you terminate the session.</p>
OID_clause	<p>lets you specify whether the object identifier (OID) of the object table should be system generated or should be based on the primary key of the table. The default is SYSTEM GENERATED.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify OBJECT IDENTIFIER IS PRIMARY KEY unless you have already specified a PRIMARY KEY constraint for the table.</li> <li>■ You cannot specify this clause for a nested table.</li> </ul> <hr/> <p><b>Note:</b> A primary key OID is locally (but not necessarily globally) unique. If you require a globally unique identifier, you must ensure that the primary key is globally unique.</p>

*OID\_index\_clause* This clause is relevant only if you have specified the *OID\_clause* as `SYSTEM GENERATED`. It specifies an index, and optionally its storage characteristics, on the hidden object identifier column.

*index* is the name of the index on the hidden system-generated object identifier column. If not specified, Oracle generates a name.

*hash\_partitioning\_clause* specifies that the table is to be partitioned using the hash method. Oracle assigns rows to partitions using a hash function on values found in columns designated as the partitioning key.

**See Also:** *Oracle8i Concepts* for more information on hash partitioning.

*column\_list* is an ordered list of columns used to determine into which partition a row belongs (**the partitioning key**).

**Restrictions:**

- You cannot specify more than 16 columns in *column\_list*.
- The *column\_list* cannot contain the `ROWID` or `UROWID` pseudocolumns.
- The columns in *column\_list* can be of any built-in datatype except `ROWID`, `LONG`, or `LOB`.

---

**Note:** If your enterprise has or will have databases using different character sets, use caution when partitioning on character columns. The sort sequence of characters is not identical in all character sets.

**See Also:** *Oracle8i National Language Support Guide* for more information on character set support.

---

You can specify hash partitioning in one of two ways:

- You can specify the number of partitions. In this case, Oracle assigns partition names of the form `SYS_Pnnn`. The `STORE IN` clause specifies one or more tablespaces where the hash partitions are to be stored. The number of tablespaces does not have to equal the number of partitions. If the number of partitions is greater than the number of tablespaces, Oracle cycles through the names of the tablespaces.
- Alternatively, you can specify individual partitions by name. The `TABLESPACE` clause specifies where the partition should be stored.

---

**Note:** The only attribute you can specify for hash partitions (or subpartitions) is `TABLESPACE`. Hash partitions inherit all other attributes from table-level defaults. Hash subpartitions inherit any attributes specified at the partition level, and inherit all other attributes from the table-level defaults.

Tablespace storage specified at the table level is overridden by tablespace storage specified at the partition level, which in turn is overridden by tablespace storage specified at the subpartition level.

---

<i>range_partitioning_clause</i>	PARTITION BY RANGE	specifies that the table is partitioned on ranges of values from <i>column_list</i> . For an index-organized table, <i>column_list</i> must be a subset of the primary key columns of the table.
	<i>column_list</i>	is an ordered list of columns used to determine into which partition a row belongs (the <b>partitioning key</b> ).  <b>Restriction:</b> The columns in <i>column_list</i> can be of any built-in datatype except ROWID, LONG, or LOB.
<i>composite_partitioning_clause</i>		specifies that table is to be first range partitioned, and then the partitions further partitioned into hash subpartitions. This combination of range partitioning and hash subpartitioning is called <b>composite partitioning</b> .
	<i>subpartition_clause</i>	specifies that Oracle should subpartition by hash each partition in <i>table</i> . The subpartitioning <i>column_list</i> is unrelated to the partitioning key, but is subject to the same restrictions.
	SUBPARTITIONS <i>quantity</i>	specifies the default number of subpartitions in each partition of <i>table</i> , and optionally one or more tablespaces in which they are to be stored.  The default value is 1. If you do not specify the <i>subpartition_clause</i> here, Oracle will create each partition with one hash subpartition unless you subsequently specify the <i>partition_level_hash_subpartitioning</i> clause.
<i>partition_definition</i>	PARTITION <i>partition</i>	specifies the physical partition attributes. If <i>partition</i> is omitted, Oracle generates a name with the form SYS_P <i>n</i> for the partition. The <i>partition</i> must conform to the rules for naming schema objects and their part as described in " <a href="#">Schema Object Naming Rules</a> " on page 2-72.

**Notes:**

- You can specify up to 64K-1 partitions and 64K-1 subpartitions. For a discussion of factors that might impose practical limits less than this number, please refer to *Oracle8i Administrator's Guide*.
- You can create a partitioned table with just one partition. Note, however, that a partitioned table with one partition is different from a nonpartitioned table. For instance, you cannot add a partition to a nonpartitioned table.

---

VALUES LESS THAN	<i>value_list</i>	specifies the noninclusive upper bound for the current partition.  is an ordered list of literal values corresponding to <i>column_list</i> in the <i>partition_by_range_clause</i> . You can substitute the keyword MAXVALUE for any literal in <i>value_list</i> . MAXVALUE specifies a maximum value that will always sort higher than any other value, including NULL.  Specifying a value other than MAXVALUE for the highest partition bound imposes an implicit integrity constraint on the table. See <i>Oracle8i Concepts</i> for more information about partition bounds.
------------------	-------------------	---

---

	<p><b>Note:</b> If <i>table</i> is partitioned on a <code>DATE</code> column, and if the NLS date format does not specify the century with the year, you must use the <code>TO_DATE</code> function with a 4-character format mask for the year. The NLS date format is determined implicitly by <code>NLS_TERRITORY</code> or explicitly by <code>NLS_DATE_FORMAT</code>.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ <i>Oracle8i National Language Support Guide</i> for more information on these initialization parameters.</li> <li>■ <a href="#">"Partitioned Table Example"</a> on page 4-414.</li> </ul>
<i>LOB_storage_clause</i>	lets you specify LOB storage characteristics for one or more LOB items in this partition. If you do not specify the <i>LOB_storage_clause</i> for a LOB item, Oracle generates a name for each LOB data partition. The system-generated names for LOB data and LOB index partitions take the form <code>SYS_LOB_Pn</code> and <code>SYS_IL_Pn</code> , respectively, where <code>P</code> stands for "partition" and <code>n</code> is a system-generated number.
<i>varray_storage_clause</i>	lets you specify storage characteristics for one or more varray items in this partition.
<i>partition_level_subpartitioning</i>	lets you specify hash subpartitions for <i>partition</i> . This clause overrides the default settings established in the <i>subpartition_clause</i> . <p><b>Restriction:</b> You can specify this clause only for a composite-partitioned table.</p> <ul style="list-style-type: none"> <li>■ You can specify individual subpartitions by name, and optionally the tablespace where each should be stored, or</li> <li>■ You can specify the number of subpartitions (and optionally one or more tablespaces where they are to be stored). In this case, Oracle assigns subpartition names of the form <code>SYS_SUBPnnn</code>. The number of tablespaces does not have to equal the number of subpartitions. If the number of partitions is greater than the number of tablespaces, Oracle cycles through the names of the tablespaces.</li> </ul>
<i>row_movement_clause</i>	determines whether a row can be moved to a different partition or subpartition because of a change to one or more of its key values during an update operation. <p><b>Restriction:</b> You can specify this clause only for a partitioned table.</p>
ENABLE	allows Oracle to move a row to a different partition or subpartition as the result of an update to the partitioning or subpartitioning key. <p><b>WARNING: Moving a row in the course of an UPDATE operation changes that row's ROWID.</b></p>
DISABLE	returns an error if an update to a partitioning or subpartitioning key would result in a row moving to a different partition or subpartition. This is the default.



---

<i>parallel_clause</i>	causes creation of the table to be parallelized, and sets the default degree of parallelism for queries and DML on the table after creation.
NOPARALLEL	specifies serial execution. This is the default.
PARALLEL	causes Oracle to select a degree of parallelism equal to the number of CPUs available on all participating instances times the value of the PARALLEL_THREADS_PER_CPU initialization parameter.
PARALLEL <i>integer</i>	specifies the <b>degree of parallelism</b> , which is the number of parallel threads used in the parallel operation. Each parallel thread may use one or two parallel execution servers. Normally, Oracle calculates the optimum degree of parallelism, so it is not necessary for you to specify <i>integer</i> .

**Restriction:** If *table* contains any columns of LOB or user-defined object type, this statement as well as subsequent INSERT, UPDATE, or DELETE operations on *table* are executed serially without notification. Subsequent queries, however, will be executed in parallel.

#### Notes

- This syntax supersedes syntax appearing in earlier releases of Oracle. Superseded syntax is still supported for backward compatibility, but may result in slightly different behavior.

**See Also:** *Oracle8i Migration*.

- A parallel hint overrides the effect of the *parallel\_clause*.
- DML statements and CREATE TABLE ... AS SELECT statements that reference remote objects can run in parallel. However, the "remote object" must really be on a remote database. The reference cannot loop back to an object on the local database (for example, by way of a synonym on the remote database pointing back to an object on the local database).

**See Also:** *Oracle8i Designing and Tuning for Performance*, *Oracle8i Concepts*, and *Oracle8i Parallel Server Concepts* for more information on parallelized operations.

<i>enable_disable_clause</i>	lets you specify whether Oracle should apply a constraint. By default, constraints are created in ENABLE VALIDATE state.
------------------------------	--

#### Restrictions:

- To enable or disable any integrity constraint, you must have defined the constraint in this or a previous statement.
- You cannot enable a referential integrity constraint unless the referenced unique or primary key constraint is already enabled.

**See Also:** "[constraint\\_clause](#)" on page 7-233 for more information on constraints.

---

**ENABLE**

specifies that the constraint will be applied to all new data in the table.

- **VALIDATE** additionally specifies that all old data also complies with the constraint. An enabled validated constraint guarantees that all data is and will continue to be valid.

If any row in the table violates the integrity constraint, the constraint remains disabled and Oracle returns an error. If all rows comply with the constraint, Oracle enables the constraint. Subsequently, if new data violates the constraint, Oracle does not execute the statement and returns an error indicating the integrity constraint violation.

If you place a primary key constraint in **ENABLE VALIDATE** mode, the validation process will verify that the primary key columns contain no nulls. To avoid this overhead, mark each column in the primary key **NOT NULL** before enabling the table's primary key constraint. (For optimal results, do this before entering data into the column.)

- **NOVALIDATE** ensures that all new DML operations on the constrained data comply with the constraint. This clause does not ensure that existing data in the table complies with the constraint and therefore does not require a table lock.
- If you specify neither **VALIDATE** nor **NOVALIDATE**, the default is **VALIDATE**.
- If you enable a unique or primary key constraint, and if no index exists on the key, Oracle creates a unique index. This index is dropped if the constraint is subsequently disabled, so Oracle rebuilds the index every time the constraint is enabled.

To avoid rebuilding the index and eliminate redundant indexes, create new primary key and unique constraints initially disabled. Then create (or use existing) nonunique indexes to enforce the constraint. Oracle does not drop a nonunique index when the constraint is disabled, so subsequent **ENABLE** operations are facilitated.

- If you change the state of any single constraint from **ENABLE NOVALIDATE** to **ENABLE VALIDATE**, the operation can be performed in parallel, and does not block reads, writes, or other DDL operations.

**Restriction:** You cannot enable a foreign key that references a unique or primary key that is disabled.

**DISABLE**

disables the integrity constraint. Disabled integrity constraints appear in the data dictionary along with enabled constraints. If you do not specify this clause when creating a constraint, Oracle automatically enables the constraint.

- **DISABLE VALIDATE** disables the constraint and drops the index on the constraint, but keeps the constraint valid. This feature is most useful in data warehousing situations, where the need arises to load into a range-partitioned table a quantity of data with a distinct range of values in the unique key. In such situations, the **disable validate** state enables you to save space by not having an index. You can then load data from a non-partitioned table into a partitioned table using the *exchange\_partition\_clause* of the **ALTER TABLE** statement or using SQL\*Loader. **All other modifications** to the table (inserts, updates, and deletes) by other SQL statements are disallowed.

If the unique key coincides with the partitioning key of the partitioned table, disabling the constraint saves overhead and has no detrimental effects. If the unique key does not coincide with the partitioning key, Oracle performs automatic table scans during the exchange to validate the constraint, which might offset the benefit of loading without an index.

- `DISABLE NOVALIDATE` signifies that Oracle makes no effort to maintain the constraint (because it is disabled) and cannot guarantee that the constraint is true (because it is not being validated). For information on when to use this setting, see *Oracle8i Designing and Tuning for Performance*.

You cannot drop a table whose primary key is being referenced by a foreign key even if the foreign key constraint is in `DISABLE NOVALIDATE` state. Further, the optimizer can use constraints in `DISABLE NOVALIDATE` state.

- If you specify neither `VALIDATE` nor `NOVALIDATE`, the default is `NOVALIDATE`.
- If you disable a unique or primary key constraint that is using a unique index, Oracle drops the unique index.

<code>UNIQUE</code>	enables or disables the unique constraint defined on the specified column or combination of columns.
<code>PRIMARY KEY</code>	enables or disables the table's primary key constraint.
<code>CONSTRAINT</code>	enables or disables the integrity constraint named <i>constraint</i> .
<i>using_index_clause</i>	<p>specifies parameters for the index Oracle creates to enforce a unique or primary key constraint. Oracle gives the index the same name as the constraint.</p> <p>You can choose the values of the <code>INITRANS</code>, <code>MAXTRANS</code>, <code>TABLESPACE</code>, <code>STORAGE</code>, and <code>PCTFREE</code> parameters for the index. These parameters are described earlier in this statement. If <i>table</i> is partitioned, you can specify a locally or globally partitioned index for the unique or primary key constraint. For a description of <code>LOCAL</code> and the <i>global_index_clause</i>, and for a description of <code>NOSORT</code> and <code>LOGGING   NOLOGGING</code> in relation to indexes, see "<a href="#">CREATE INDEX</a>" on page 7-291.</p> <p><b>Restriction:</b> Use these parameters only when enabling unique and primary key constraints.</p>
<code>EXCEPTIONS INTO</code>	specifies a table into which Oracle places the rowids of all rows violating the constraint. If you omit schema, Oracle assumes the exceptions table is in your own schema. If you omit this clause altogether, Oracle assumes that the table is named <code>EXCEPTIONS</code> . The exceptions table must be on your local database.

You can create the `EXCEPTIONS` table using one of these scripts:

- `UTLEXCPT.SQL` uses physical rowids. Therefore it can accommodate rows from conventional tables but not from index-organized tables. (See the Note that follows.)
- `UTLEXPT1.SQL` uses universal rowids, so it can accommodate rows from both conventional and index-organized tables.

If you create your own exceptions table, it must follow the format prescribed by one of these two scripts. See *Oracle8i Migration* for compatibility issues related to the use of these scripts.

---

**Note:** If you are collecting exceptions from index-organized tables based on primary keys (rather than universal rowids), you must create a separate exceptions table for each index-organized table to accommodate its primary-key storage. You create multiple exceptions tables with different names by modifying and resubmitting the script.

For information on the SQL scripts, see the `DBMS_IOT` package in *Oracle8i Supplied PL/SQL Packages Reference*. For information on eliminating migrated and chained rows, see *Oracle8i Designing and Tuning for Performance*.

---

CASCADE

disables any integrity constraints that depend on the specified integrity constraint. To disable a primary or unique key that is part of a referential integrity constraint, you must specify this clause.

**Restriction:** You can specify `CASCADE` only if you have specified `DISABLE`.

CACHE

for data that will be accessed frequently, specifies that the blocks retrieved for this table are placed at the most recently used end of the LRU list in the buffer cache when a full table scan is performed. This clause is useful for small lookup tables.

As a parameter in the *LOB\_storage\_clause*, `CACHE` specifies that Oracle places the LOB values in the buffer cache for faster access.

**Restriction:** You cannot specify `CACHE` for an index-organized table. However, index-organized tables implicitly provide `CACHE` behavior.

NOCACHE

for data that will not be accessed frequently, specifies that the blocks retrieved for this table are placed at the least recently used end of the LRU list in the buffer cache when a full table scan is performed. This is the default.

As a parameter in the *LOB\_storage\_clause*, `NOCACHE` specifies that the LOB value either is not brought into the buffer cache or is brought into the buffer cache and placed at the least recently used end of the LRU list. (The latter is the default behavior.)

**Restriction:** You cannot specify `NOCACHE` for an index-organized table.

CACHE READS

applies only to LOB storage. It specifies that LOB values are brought into the buffer cache only during read operations, but not during write operations.

MONITORING

specifies that modification statistics can be collected on this table. These statistics are estimates of the number of rows affected by DML statements over a particular period of time. They are available for use by the optimizer or for analysis by the user.

**Restriction:** You cannot specify `MONITORING` for a temporary table.

---

NOMONITORING	<p>specifies that the table will not have modification statistics collected. This is the default.</p> <p><b>Restriction:</b> You cannot specify <code>NOMONITORING</code> for a temporary table.</p>
<code>AS subquery</code>	<p>inserts the rows returned by the subquery into the table upon its creation. See "<a href="#">SELECT and Subqueries</a>" on page 7-569.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ The number of columns in the table must equal the number of expressions in the subquery.</li> <li>■ The column definitions can specify only column names, default values, and integrity constraints, not datatypes.</li> <li>■ You cannot define a referential integrity constraint in a <code>CREATE TABLE</code> statement that contains <code>AS subquery</code>. Instead, you must create the table without the constraint and then add it later with an <code>ALTER TABLE</code> statement.</li> </ul> <p>If you specify the <i>parallel_clause</i> in this statement, Oracle will ignore any value you specify for the <code>INITIAL</code> storage parameter, and will instead use the value of the <code>NEXT</code> parameter. For information on these parameters, see the "<a href="#">storage_clause</a>" on page 7-605.</p> <p>Oracle derives datatypes and lengths from the subquery. Oracle also follows the following rules for integrity constraints:</p> <ul style="list-style-type: none"> <li>■ Oracle automatically defines any <code>NOT NULL</code> constraints on columns in the new table that existed on the corresponding columns of the selected table if the subquery selects the column rather than an expression containing the column.</li> <li>■ If a <code>CREATE TABLE</code> statement contains both <code>AS subquery</code> and a <code>CONSTRAINT</code> clause or an <code>ENABLE</code> clause with the <code>EXCEPTIONS INTO</code> clause, Oracle ignores <code>AS subquery</code>. If any rows violate the constraint, Oracle does not create the table and returns an error.</li> </ul> <p>If all expressions in <i>subquery</i> are columns, rather than expressions, you can omit the columns from the table definition entirely. In this case, the names of the columns of table are the same as the columns in <i>subquery</i>.</p> <p>You can use <i>subquery</i> in combination with the <code>TO_LOB</code> function to convert the values in a <code>LONG</code> column in another table to LOB values in a column of the table you are creating. For a discussion of why and when to copy <code>LONGs</code> to LOBs, see <i>Oracle8i Migration</i>. For a description of how to use the <code>TO_LOB</code> function, see "<a href="#">Conversion Functions</a>" on page 4-5.</p> <hr/> <p><b>Note:</b> If <i>subquery</i> returns (in part or totally) the equivalent of an existing materialized view, Oracle may use the materialized view (for query rewrite) in place of one or more tables specified in <i>subquery</i>.</p> <p><b>See also:</b> <i>Oracle8i Data Warehousing Guide</i> for more information on materialized views and query rewrite.</p> <hr/> <p><i>order_by_clause</i> orders rows returned by the statements.</p> <p><b>See also:</b> "<a href="#">SELECT and Subqueries</a>" on page 7-569 for more information on the <i>order_by_clause</i>.</p>

---

**Note:** When specified with `CREATE TABLE`, this clause does not necessarily order data cross the entire table. (For example, it does not order across partitions.) Specify this clause if you intend to create an index on the same key as the `ORDER BY` key column. Oracle will cluster data on the `ORDER BY` key so that it corresponds to the index key.

---

For object tables, *subquery* can contain either one expression corresponding to the table type, or the number of top-level attributes of the table type.

---

## Examples

**General Example** To define the `EMP` table owned by `SCOTT`, you could issue the following statement:

```
CREATE TABLE scott.emp
  (empno      NUMBER          CONSTRAINT pk_emp PRIMARY KEY,
   ename      VARCHAR2(10)   CONSTRAINT nn_ename NOT NULL
                                CONSTRAINT upper_ename
                                CHECK (ename = UPPER(ename)),
   job        VARCHAR2(9),
   mgr        NUMBER          CONSTRAINT fk_mgr
                                REFERENCES scott.emp(empno),
   hiredate   DATE            DEFAULT SYSDATE,
   sal        NUMBER(10,2)   CONSTRAINT ck_sal
                                CHECK (sal > 500),
   comm       NUMBER(9,0)     DEFAULT NULL,
   deptno     NUMBER(2)      CONSTRAINT nn_deptno NOT NULL
                                CONSTRAINT fk_deptno
                                REFERENCES scott.dept(deptno) )

PCTFREE 5 PCTUSED 75;
```

This table contains eight columns. The `EMPNO` column is of datatype `NUMBER` and has an associated integrity constraint named `PK_EMP`. The `HIRDEDATE` column is of datatype `DATE` and has a default value of `SYSDATE`, and so on.

This table definition specifies a `PCTFREE` of 5 and a `PCTUSED` of 75, which is appropriate for a relatively static table. The definition also defines integrity constraints on some columns of the `EMP` table.

**Temporary Table Example** The following statement creates a temporary table `FLIGHT_SCHEDULE` for use in an automated airline reservation scheduling system. Each client has its own session and can store temporary schedules. The temporary schedules are deleted at the end of the session.

```
CREATE GLOBAL TEMPORARY TABLE flight_schedule (
  startdate DATE,
  enddate DATE,
  cost NUMBER)
ON COMMIT PRESERVE ROWS;
```

**Storage Example** To define the sample table SALGRADE in the HUMAN\_RESOURCE tablespace with a small storage capacity and limited allocation potential, issue the following statement:

```
CREATE TABLE salgrade
  ( grade NUMBER CONSTRAINT pk_salgrade
    PRIMARY KEY
    USING INDEX TABLESPACE users_a,
    losal NUMBER,
    hisal NUMBER )
TABLESPACE human_resource
STORAGE (INITIAL      6144
         NEXT         6144
         MINEXTENTS   1
         MAXEXTENTS   5 );
```

The above statement also defines a primary key constraint on the GRADE column and specifies that the index Oracle creates to enforce this constraint is created in the USERS\_A tablespace.

**See Also:** The "[constraint\\_clause](#)" on page 7-233 for more examples of defining integrity constraints.

**PARALLEL Example** The following statement creates a table using an optimum number of parallel execution servers to scan SCOTT.EMP and to populate EMP\_DEPT:

```
CREATE TABLE emp_dept
  PARALLEL
  AS SELECT * FROM scott.emp
  WHERE deptno = 10;
```

Using parallelism speeds up the creation of the table because Oracle uses parallel execution servers to create the table. After the table is created, querying the table is also faster, because the same degree of parallelism is used to access the table.

**NOPARALLEL Example** The following statement creates a table serially. Subsequent DML and queries on the table will also be serially executed.

## CREATE TABLE

---

```
CREATE TABLE emp_dept
AS SELECT * FROM scott.emp
WHERE deptno = 10;
```

**ENABLE VALIDATE Example** The following statement creates the DEPT table, defines a primary key constraint, and places it in ENABLE VALIDATE state:

```
CREATE TABLE dept
(deptno NUMBER (2) PRIMARY KEY,
 dname VARCHAR2(10),
 loc VARCHAR2(9) )
TABLESPACE user_a;
```

**DISABLE Example** The following statement creates the DEPT table and defines a disabled primary key constraint:

```
CREATE TABLE dept
(deptno NUMBER (2) PRIMARY KEY DISABLE,
 dname VARCHAR2(10),
 loc VARCHAR2(9) );
```

**EXCEPTIONS INTO Example** The following example creates the ORDER\_EXCEPTIONS table to hold rows from an index-organized table ORDERS that violate integrity constraint CHECK\_ORDERS:

```
CREATE TABLE orders
(ord_num NUMBER PRIMARY KEY,
 ord_quantity NUMBER)
ORGANIZATION INDEX;

EXECUTE DBMS_IOT.BUILD_EXCEPTIONS_TABLE
('SCOTT', 'ORDERS', 'ORDER_EXCEPTIONS');

ALTER TABLE orders
ADD CONSTRAINT CHECK_ORDERS CHECK (ord_quantity > 0)
EXCEPTIONS INTO ORDER_EXCEPTIONS;
```

To specify an exception table, you must have the privileges necessary to insert rows into the table. To examine the identified exceptions, you must have the privileges necessary to query the exceptions table.



**See Also:**

- ["INSERT"](#) on page 7-539.
- ["SELECT and Subqueries"](#) on page 7-569 for information on the privileges necessary to insert rows into tables.

**Nested Table Example** The following statement creates relational table `EMPLOYEE` with a nested table column `PROJECTS`:

```
CREATE TABLE employee
  (empno NUMBER, name CHAR(31), projects PROJ_TABLE_TYPE)
  NESTED TABLE projects STORE AS nested_proj_table(
    (PRIMARY KEY (nested_table_id, pno)) ORGANIZATION INDEX)
  RETURN AS LOCATOR;
```

**LOB Column Example** The following statement creates table `LOB_TAB` with two LOB columns and specifies the LOB storage characteristics:

```
CREATE TABLE lob_tab (col1 BLOB, col2 CLOB)
  STORAGE (INITIAL 256 NEXT 256)
  LOB (col1, col2) STORE AS
    (TABLESPACE lob_seg_ts
     STORAGE (INITIAL 6144 NEXT 6144)
     CHUNK 4000
     NOCACHE LOGGING);
```

In the example, Oracle rounds the `CHUNK` up to 4096 (the nearest multiple of the block size of 2048).

**Index-Organized Table Example** The following statement creates an index-organized table:

```
CREATE TABLE docindex
  ( token          CHAR(20),
    doc_oid        INTEGER,
    token_frequency SMALLINT,
    token_occurrence_data VARCHAR2(512),
    CONSTRAINT pk_docindex PRIMARY KEY (token, doc_oid) )
  ORGANIZATION INDEX TABLESPACE text_collection
  PCTTHRESHOLD 20 INCLUDING token_frequency
  OVERFLOW TABLESPACE text_collection_overflow;
```

**Partitioned Table Example** The following statement creates a table with three partitions:

```
CREATE TABLE stock_xactions
  (stock_symbol CHAR(5),
   stock_series CHAR(1),
   num_shares NUMBER(10),
   price NUMBER(5,2),
   trade_date DATE)
STORAGE (INITIAL 100K NEXT 50K) LOGGING
PARTITION BY RANGE (trade_date)
  (PARTITION sx1992 VALUES LESS THAN (TO_DATE('01-JAN-1993', 'DD-MON-YYYY'))
   TABLESPACE ts0 NOLOGGING,
   PARTITION sx1993 VALUES LESS THAN (TO_DATE('01-JAN-1994', 'DD-MON-YYYY'))
   TABLESPACE ts1,
   PARTITION sx1994 VALUES LESS THAN (TO_DATE('01-JAN-1995', 'DD-MON-YYYY'))
   TABLESPACE ts2);
```

For information about partitioned table maintenance operations, see the *Oracle8i Administrator's Guide*.

**Partitioned Table with LOB Columns Example** This statement creates a partitioned table PT with two partitions P1 and P2, and three LOB columns, B, C, and D:

```
CREATE TABLE PT (A NUMBER, B BLOB, C CLOB, D CLOB)
  LOB (B,C,D) STORE AS (STORAGE (NEXT 20M))
  PARTITION BY RANGE (A)
  (PARTITION P1 VALUES LESS THAN (10) TABLESPACE TS1
   LOB (B,D) STORE AS (TABLESPACE TSA STORAGE (INITIAL 20M)),
   PARTITION P2 VALUES LESS THAN (20)
   LOB (B,C) STORE AS (TABLESPACE TSB)
  TABLESPACE TSX;
```

Partition P1 will be in tablespace TS1. The LOB data partitions for B and D will be in tablespace TSA. The LOB data partition for C will be in tablespace TS1. The storage attribute INITIAL is specified for LOB columns B and D; other attributes will be inherited from the default table-level specification. The default LOB storage attributes not specified at the table level will be inherited from the tablespace TSA for columns B and D and tablespace TS1 for column C. LOB index partitions will be in the same tablespaces as the corresponding LOB data partitions. Other storage attributes will be based on values of the corresponding attributes of the LOB data partitions and default attributes of the tablespace where the index partitions reside.

Partition P2 will be in the default tablespace TSX. The LOB data for B and C will be in tablespace TSB. The LOB data for D will be in tablespace TSX. The LOB index for columns B and C will be in tablespace TSB. The LOB index for column D will be in tablespace TSX.

**Hash-Partitioned Table Example** This statement creates a table partitioned by hash on columns containing data about chemicals. The hash partitions are stored in tablespaces TBS1, TBS2, TBS3, and TBS4:

```
CREATE TABLE exp_data (
    d DATE, temperature NUMBER, Fe2O3_concentration NUMBER,
    HCl_concentration NUMBER, Au_concentration NUMBER,
    amps NUMBER, observation VARCHAR(4000))
PARTITION BY HASH (HCl_concentration, Au_concentration)
PARTITIONS 32 STORE IN (tbs1, tbs2, tbs3, tbs4);
```

**Composite-Partitioned Table Example** This statement creates a composite-partitioned table. The range partitioning facilitates data and partition pruning by sale date. The hash subpartitioning enables subpartition elimination for queries by a specific item number. Most of the partitions consist of 8 subpartitions. However, the partition covering the slowest quarter will have 4 subpartitions, and the partition covering the busiest quarter will have 16 subpartitions.

```
CREATE TABLE sales (item INTEGER, qty INTEGER,
    store VARCHAR(30),
    dept NUMBER, sale_date DATE)
PARTITION BY RANGE (sale_date)
SUBPARTITION BY HASH(item)
SUBPARTITIONS 8
STORE IN (tbs1, tbs2, tbs3, tbs4, tbs5, tbs6, tbs7, tbs8)
(PARTITION q1_1997
    VALUES LESS THAN (TO_DATE('01-apr-1997', 'dd-mon-yyyy')),
PARTITION q2_1997
    VALUES LESS THAN (TO_DATE('01-jul-1997', 'dd-mon-yyyy')),
PARTITION q3_1997
    VALUES LESS THAN (TO_DATE('01-oct-1997', 'dd-mon-yyyy'))
    (SUBPARTITION q3_1997_s1 TABLESPACE ts1,
    SUBPARTITION q3_1997_s2 TABLESPACE ts3,
    SUBPARTITION q3_1997_s3 TABLESPACE ts5,
    SUBPARTITION q3_1997_s4 TABLESPACE ts7),
PARTITION q4_1997
    VALUES LESS THAN (TO_DATE('01-jan-1998', 'dd-mon-yyyy'))
SUBPARTITIONS 16
STORE IN (tbs1, tbs3, tbs5, tbs7, tbs8, tbs9, tbs10,
```

```
        tbs11),
PARTITION q1_1998
VALUES LESS THAN (TO_DATE('01-apr-1998', 'dd-mon-yyyy')));
```

**Object Table Examples** Consider object type DEPT\_T:

```
CREATE TYPE dept_t AS OBJECT
(  dname VARCHAR2(100),
  address VARCHAR2(200) );
```

Object table DEPT holds department objects of type DEPT\_T:

```
CREATE TABLE dept OF dept_t;
```

The following statement creates object table SALESREPS with a user-defined object type, SALESREP\_T:

```
CREATE OR REPLACE TYPE salesrep_t AS OBJECT
(  repId NUMBER,
  repName VARCHAR2(64));
CREATE TABLE salesreps OF salesrep_t;
```

**Nested Table Example** The following statement creates relational table EMPLOYEE with a nested table column PROJECTS:

```
CREATE TABLE employee (empno NUMBER, name CHAR(31),
  projects PROJ_TABLE_TYPE)
  NESTED TABLE projects STORE AS nested_proj_table;
```

**REF Example** The following example creates object type DEPT\_T and object table DEPT to store instances of all departments. A table with a scoped REF is then created.

```
CREATE TYPE dept_t AS OBJECT
(  dname  VARCHAR2(100),
  address VARCHAR2(200) );

CREATE TABLE dept OF dept_t;

CREATE TABLE emp
(  ename  VARCHAR2(100),
  enumber NUMBER,
  edept  REF dept_t SCOPE IS dept );
```

The following statement creates a table with a REF column which has a referential constraint defined on it:

```
CREATE TABLE emp
  ( ename   VARCHAR2(100),
    enumber NUMBER,
    edept   REF dept_t REFERENCES dept);
```

**User-Defined OID Example** This example creates an object type and a corresponding object table whose OID is primary key based:

```
CREATE TYPE emp_t AS OBJECT (empno NUMBER, address CHAR(30));
CREATE TABLE emp OF emp_t (empno PRIMARY KEY)
  OBJECT IDENTIFIER IS PRIMARY KEY;
```

You can subsequently reference the EMP object table in either of the following two ways:

```
CREATE TABLE dept (dno NUMBER
  mgr_ref REF emp_t SCOPE IS emp);
CREATE TABLE dept (
  dno NUMBER,
  mgr_ref REF emp_t CONSTRAINT mgr_in_emp REFERENCES emp);
```

**Constraints on Type Columns Example**

```
CREATE TYPE address AS OBJECT
  ( hno   NUMBER,
    street VARCHAR2(40),
    city  VARCHAR2(20),
    zip   VARCHAR2(5),
    phone VARCHAR2(10) );

CREATE TYPE person AS OBJECT
  ( name      VARCHAR2(40),
    dateofbirth DATE,
    homeaddress address,
    manager   REF person );

CREATE TABLE persons OF person
  ( homeaddress NOT NULL
    UNIQUE (homeaddress.phone),
    CHECK (homeaddress.zip IS NOT NULL),
    CHECK (homeaddress.city <> 'San Francisco') );
```

**PARALLEL Example** The following statement creates a table using 10 parallel execution servers, 5 to scan SCOTT.EMP and another 5 to populate EMP\_DEPT:

```
CREATE TABLE emp_dept
```

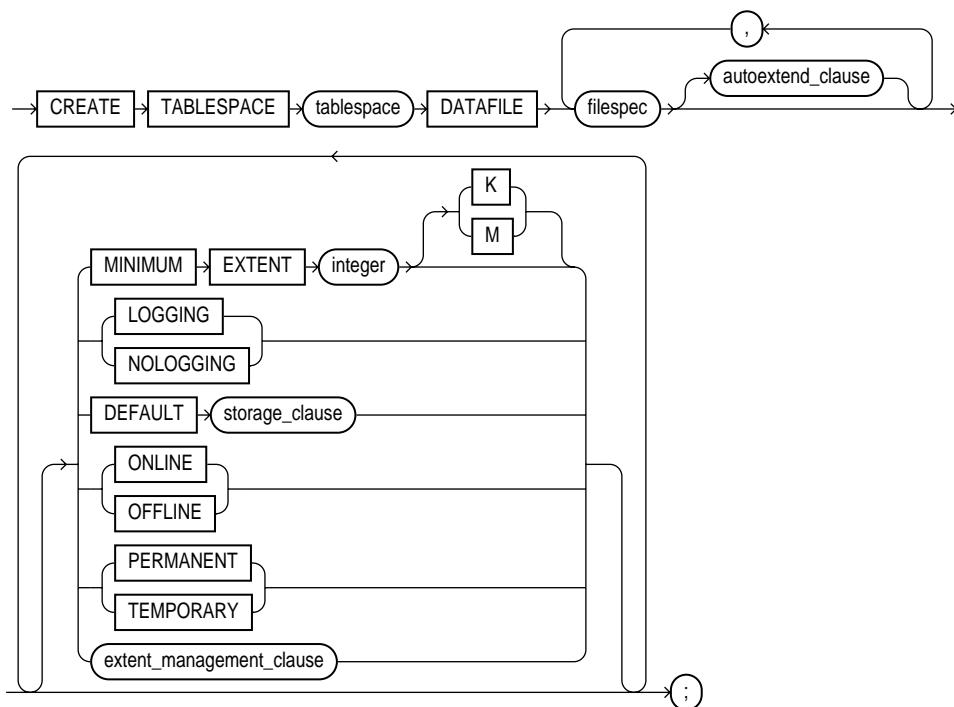
## CREATE TABLE

---

```
PARALLEL (5)  
AS SELECT * FROM scott.emp  
WHERE deptno = 10;
```


# CREATE TABLESPACE


## Syntax




**filespec:** See "filespec" on page 7-516.

**autoextend\_clause::=**



**maxsize\_clause::=**

**storage\_clause:** See "storage\_clause" on page 7-605.

**extent\_management\_clause::=****Purpose**

To create a tablespace. A **tablespace** is an allocation of space in the database that can contain schema objects. For information on tablespaces, see *Oracle8i Concepts*.

When you create a tablespace, it is initially a read-write tablespace. You can subsequently use the ALTER TABLESPACE statement to take the tablespace offline or online, add datafiles to it, or make it a read-only tablespace. See "ALTER TABLESPACE" on page 7-176.

You can also drop a tablespace from the database with the DROP TABLESPACE statement. See "DROP TABLESPACE" on page 7-503.

**Prerequisites**

You must have CREATE TABLESPACE system privilege. Also, the SYSTEM tablespace must contain at least two rollback segments including the SYSTEM rollback segment.

Before you can create a tablespace, you must create a database to contain it, and the database must be open. See "CREATE DATABASE" on page 7-267.



## Keywords and Parameters

---

<i>tablespace</i>	is the name of the tablespace to be created.
DATAFILE <i>filespec</i>	specifies the datafile or files to make up the tablespace. See " <a href="#">filespec</a> " on page 7-516.
	<b>Note:</b> For operating systems that support raw devices, the <i>filespec</i> REUSE keyword has no meaning when specifying a raw device as a datafile. Such a CREATE TABLESPACE statement will succeed whether or not you specify REUSE.
<i>autoextend_</i> <i>clause</i>	enables or disables the automatic extension of the datafile.
	OFF                    disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in further ALTER TABLESPACE AUTOEXTEND statements.
	ON                     enables autoextend.
	NEXT                  specifies the disk space to allocate to the datafile when more extents are required.
	<i>maxsize_clause</i> specifies the maximum disk space allowed for allocation to the datafile.
	UNLIMITED           sets no limit on allocating disk space to the datafile.
MINIMUM EXTENT <i>integer</i>	controls free space fragmentation in the tablespace by ensuring that every used or free extent size in a tablespace is at least as large as, and is a multiple of, <i>integer</i> . <b>See also:</b> <i>Oracle8i Concepts</i> for more information about using MINIMUM EXTENT to control fragmentation.
	<b>Note:</b> This clause is not relevant for a dictionary-managed temporary tablespace.
LOGGING   NOLOGGING	specifies the default logging attributes of all tables, indexes, and partitions within the tablespace. LOGGING is the default.  The tablespace-level logging attribute can be overridden by logging specifications at the table, index, and partition levels.  Only the following operations support the NOLOGGING mode: <b>DML:</b> direct-load INSERT (serial or parallel), Direct Loader (SQL*Loader) <b>DDL:</b> CREATE TABLE ... AS SELECT, CREATE INDEX, ALTER INDEX ... REBUILD, ALTER INDEX ... REBUILD PARTITION, ALTER INDEX ... SPLIT PARTITION, ALTER TABLE ... SPLIT PARTITION, and ALTER TABLE ... MOVE PARTITION

	<p>In <code>NOLOGGING</code> mode, data is modified with minimal logging (to mark new extents <code>INVALID</code> and to record dictionary changes). When applied during media recovery, the extent invalidation records mark a range of blocks as logically corrupt, because the redo data is not logged. Therefore, if you cannot afford to lose the object, you should take a backup after the <code>NOLOGGING</code> operation.</p>
<code>DEFAULT storage_clause</code>	<p>specifies the default storage parameters for all objects created in the tablespace. For a dictionary-managed temporary tablespace, Oracle considers only the <code>NEXT</code> parameter of the <code>storage_clause</code>. For information on storage parameters, see the "<a href="#">storage_clause</a>" on page 7-605.</p>
<code>ONLINE</code>	<p>makes the tablespace available immediately after creation to users who have been granted access to the tablespace. This is the default.</p>
<code>OFFLINE</code>	<p>makes the tablespace unavailable immediately after creation.</p> <p>The data dictionary view <code>DBA_TABLESPACES</code> indicates whether each tablespace is online or offline.</p>
<code>PERMANENT</code>	<p>specifies that the tablespace will be used to hold permanent objects. This is the default.</p>
<code>TEMPORARY</code>	<p>specifies that the tablespace will be used only to hold temporary objects, for example, segments used by implicit sorts to handle <code>ORDER BY</code> clauses.</p> <p><b>Restriction:</b> If you specify <code>TEMPORARY</code>, you cannot specify <code>EXTENT MANAGEMENT LOCAL</code>.</p>
<code>extent_management_clause</code>	<p>specifies how the extents of the tablespace will be managed.</p>
<code>DICTIONARY</code>	<p>specifies that the tablespace is managed using dictionary tables. This is the default.</p>
<code>LOCAL</code>	<p>specifies that tablespace is locally managed. Locally managed tablespaces have some part of the tablespace set aside for a bitmap. For a discussion of locally managed tablespaces, see <i>Oracle8i Concepts</i>.</p> <ul style="list-style-type: none"><li>■ <code>AUTOALLOCATE</code> specifies that the tablespace is system managed. Users cannot specify an extent size.</li><li>■ <code>UNIFORM</code> specifies that the tablespace is managed with uniform extents of <code>SIZE</code> bytes. Use <code>K</code> or <code>M</code> to specify the extent size in kilobytes or megabytes. The default <code>SIZE</code> is 1 megabyte.</li></ul> <p>If you do not specify either <code>AUTOALLOCATE</code> or <code>UNIFORM</code>, then <code>AUTOALLOCATE</code> is the default.</p> <p><b>Restriction:</b> If you specify <code>LOCAL</code>, you cannot specify <code>DEFAULT storage_clause</code>, <code>MINIMUM EXTENT</code>, or <code>TEMPORARY</code>.</p>

---

## Examples

**DEFAULT Storage Example** This statement creates a tablespace named `TABSPACE_2` with one datafile:

```
CREATE TABLESPACE tabspace_2
  DATAFILE 'diska:tabspace_file2.dat' SIZE 20M
  DEFAULT STORAGE (INITIAL 10K NEXT 50K
                  MINEXTENTS 1 MAXEXTENTS 999)
  ONLINE;
```

**AUTOEXTEND Example** This statement creates a tablespace named `TABSPACE_3` with one datafile. When more space is required, 50 kilobyte extents will be added up to a maximum size of 10 megabytes:

```
CREATE TABLESPACE tabspace_5
  DATAFILE 'diskb:tabspace_file3.dat' SIZE 500K REUSE
  AUTOEXTEND ON NEXT 500K MAXSIZE 10M;
```

**MINIMUM EXTENT Example** This statement creates tablespace `TABSPACE_5` with one datafile and allocates every extent as a multiple of 64K:

```
CREATE TABLESPACE tabspace_3
  DATAFILE 'tabspace_file5.dbf' SIZE 2M
  MINIMUM EXTENT 64K
  DEFAULT STORAGE (INITIAL 128K NEXT 128K)
  LOGGING;
```


**Locally Managed Example** In the following statement, we assume that the database block size is 2K.

```
CREATE TABLESPACE tbs_1 DATAFILE 'file_1.f' SIZE 10M
  EXTENT MANAGEMENT LOCAL UNIFORM SIZE 128K;
```

This statement creates a locally managed tablespace in which every extent is 128K and each bit in the bit map describes 64 blocks.


# CREATE TEMPORARY TABLESPACE

## Syntax




**filespec:** See "filespec" on page 7-516.

**autoextend\_clause::=**



**maxsize\_clause::=**



## Purpose

To create a **temporary tablespace**, which is an allocation of space in the database that can contain schema objects for the duration of a session.

## Prerequisites

You must have the `CREATE TABLESPACE` system privilege.

## Keywords and Parameters

<i>tablespace</i>	is the name of the temporary tablespace.												
TEMPFILE <i>filespec</i>	specifies the tempfiles that make up the tablespace. See " <a href="#">filespec</a> " on page 7-516. <b>Note:</b> Media recovery does not recognize tempfiles.												
<i>autoextend_clause</i>	enables or disables the automatic extension of the tempfile. <table border="0"> <tr> <td>OFF</td> <td>disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in further ALTER TABLESPACE AUTOEXTEND statements.</td> </tr> <tr> <td>ON</td> <td>enables autoextend.</td> </tr> <tr> <td>NEXT</td> <td>specifies the disk space to allocate to the tempfile when more extents are required.</td> </tr> <tr> <td><i>maxsize_clause</i></td> <td>specifies the maximum disk space allowed for allocation to the tempfile.</td> </tr> <tr> <td><i>integer</i></td> <td>specifies in bytes the maximum disk space allowed for allocation to the tempfile. Use K or M to specify this space in kilobytes or megabytes.</td> </tr> <tr> <td>UNLIMITED</td> <td>sets no limit on allocating disk space to the tempfile.</td> </tr> </table>	OFF	disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in further ALTER TABLESPACE AUTOEXTEND statements.	ON	enables autoextend.	NEXT	specifies the disk space to allocate to the tempfile when more extents are required.	<i>maxsize_clause</i>	specifies the maximum disk space allowed for allocation to the tempfile.	<i>integer</i>	specifies in bytes the maximum disk space allowed for allocation to the tempfile. Use K or M to specify this space in kilobytes or megabytes.	UNLIMITED	sets no limit on allocating disk space to the tempfile.
OFF	disables autoextend if it is turned on. NEXT and MAXSIZE are set to zero. Values for NEXT and MAXSIZE must be respecified in further ALTER TABLESPACE AUTOEXTEND statements.												
ON	enables autoextend.												
NEXT	specifies the disk space to allocate to the tempfile when more extents are required.												
<i>maxsize_clause</i>	specifies the maximum disk space allowed for allocation to the tempfile.												
<i>integer</i>	specifies in bytes the maximum disk space allowed for allocation to the tempfile. Use K or M to specify this space in kilobytes or megabytes.												
UNLIMITED	sets no limit on allocating disk space to the tempfile.												
EXTENT MANAGEMENT LOCAL	specifies that the tablespace is locally managed, meaning that some part of the tablespace is set aside for a bitmap. For a discussion of locally managed tablespaces, see <i>Oracle8i Concepts</i> . <table border="0"> <tr> <td>UNIFORM</td> <td>determines the size of the extents of the temporary tablespace in bytes. All extents of temporary tablespaces are the same size (uniform). If you do not specify this clause, Oracle uses uniform extents of 1M.</td> </tr> <tr> <td>SIZE</td> <td>specifies in bytes the size of the tablespace extents. Use K or M to specify the size in kilobytes or megabytes. If you do not specify SIZE, Oracle uses the default extent size of 1M.</td> </tr> </table>	UNIFORM	determines the size of the extents of the temporary tablespace in bytes. All extents of temporary tablespaces are the same size (uniform). If you do not specify this clause, Oracle uses uniform extents of 1M.	SIZE	specifies in bytes the size of the tablespace extents. Use K or M to specify the size in kilobytes or megabytes. If you do not specify SIZE, Oracle uses the default extent size of 1M.								
UNIFORM	determines the size of the extents of the temporary tablespace in bytes. All extents of temporary tablespaces are the same size (uniform). If you do not specify this clause, Oracle uses uniform extents of 1M.												
SIZE	specifies in bytes the size of the tablespace extents. Use K or M to specify the size in kilobytes or megabytes. If you do not specify SIZE, Oracle uses the default extent size of 1M.												

## Example

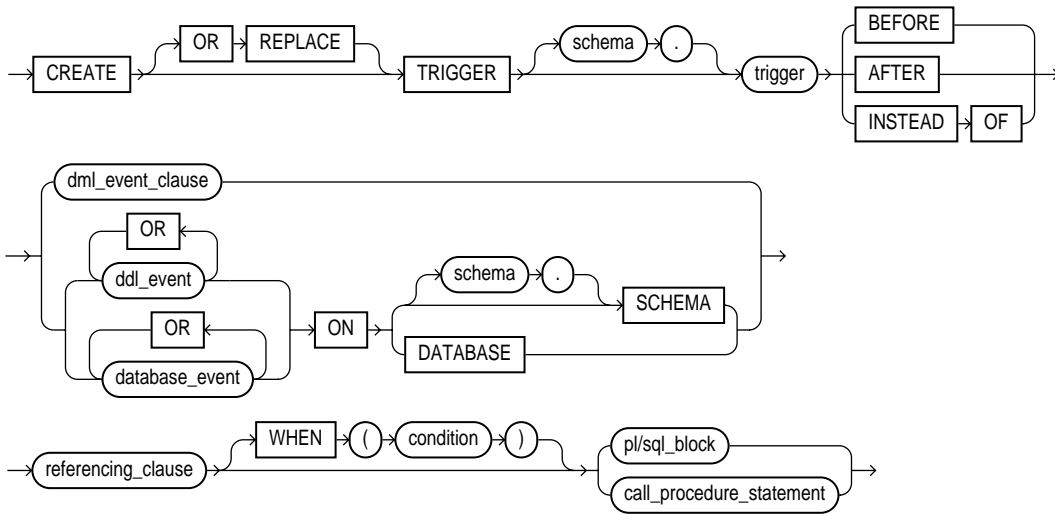
This statement creates a temporary tablespace in which each extent is 16M.

```
CREATE TEMPORARY TABLESPACE tbs_1 TEMPFILE 'file_1.f'
  EXTENT MANAGEMENT LOCAL UNIFORM SIZE 16M;
```

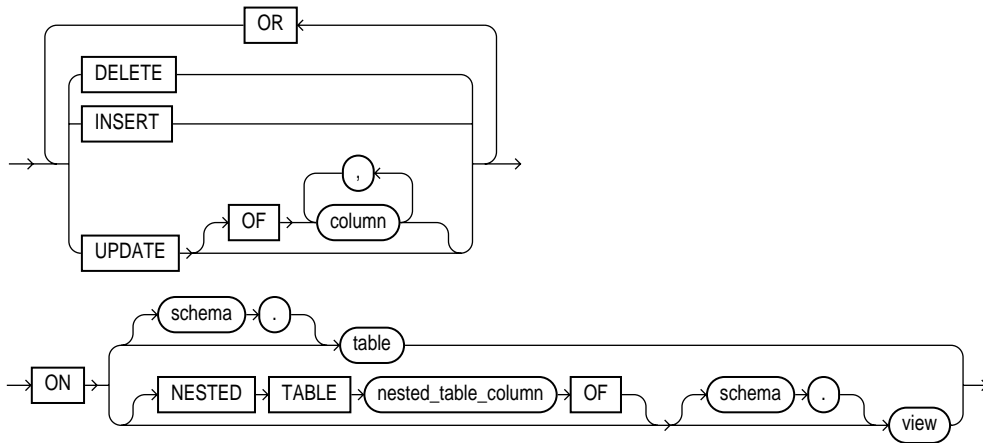
If we assume the default database block size of 2K, and that each bit in the map represents one extent, then each bit maps 8,000 blocks.

# CREATE TRIGGER

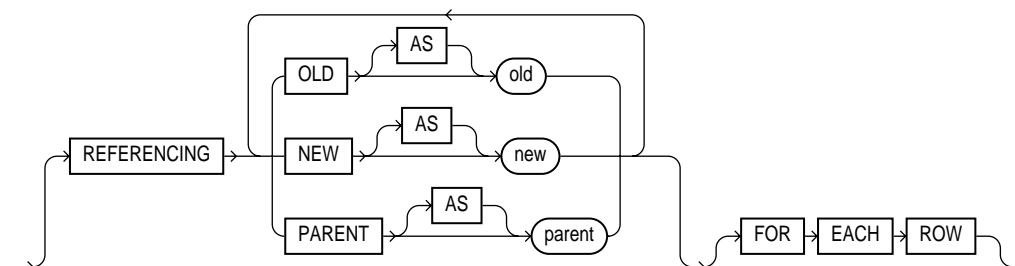
## Syntax



### dml\_event\_clause::=



`referencing_clause ::=`



## Purpose

To create and enable a database trigger. A **database trigger** is

- A stored PL/SQL block associated with a table, a schema, or the database
- An anonymous PL/SQL block or a call to a procedure implemented in PL/SQL or Java

Oracle automatically executes a trigger when specified conditions occur.

### See Also:

- *Oracle8i Concepts* for a description of the various types of triggers.
- *Oracle8i Application Developer's Guide - Fundamentals* for more information on how to design triggers for the above purposes.

## Prerequisites

Before a trigger can be created, the user SYS must run the SQL script `DBMSSTDY.SQL`. The exact name and location of this script depend on your operating system.

- To create a trigger in your own schema on a table in your own schema or on your own schema (SCHEMA), you must have the `CREATE TRIGGER` privilege.
- To create a trigger in any schema on a table in any schema, or on another user's schema (*schema*.SCHEMA), you must have the `CREATE ANY TRIGGER` privilege.
- In addition to the preceding privileges, to create a trigger on `DATABASE`, you must have the `ADMINISTER DATABASE TRIGGER` system privilege.

If the trigger issues SQL statements or calls procedures or functions, then the owner of the trigger must have the privileges necessary to perform these operations. These privileges must be granted directly to the owner, rather than acquired through roles.

## Keywords and Parameters

CREATE	<p>creates a new trigger. When you create a trigger, Oracle enables it automatically. You can subsequently disable and enable a trigger with the <code>DISABLE</code> and <code>ENABLE</code> clause of the <code>ALTER TRIGGER</code> or <code>ALTER TABLE</code> statement. For information on how to enable and disable triggers, see "<a href="#">ALTER TRIGGER</a>" on page 7-183 and "<a href="#">ALTER TABLE</a>" on page 7-123.</p> <p>If a trigger produces compilation errors, it is still created, but it fails on execution. You can see the associated compiler error messages with the SQL*Plus command <code>SHOW ERRORS</code>. This means it effectively blocks all triggering DML statements until it is disabled, replaced by a version without compilation errors, or dropped.</p>
OR REPLACE	re-creates the trigger if it already exists. Use this clause to change the definition of an existing trigger without first dropping it.
<i>schema</i>	is the schema to contain the trigger. If you omit <i>schema</i> , Oracle creates the trigger in your own schema.
<i>trigger</i>	is the name of the trigger to be created.
	<p><b>Note:</b> If you create a trigger on a base table of a materialized view, you must ensure that the trigger does not fire during a refresh of the materialized view. (During refresh, the <code>DBMS_SNAPSHOT</code> procedure <code>I_AM_A_REFRESH</code> returns <code>TRUE</code>.)</p>
BEFORE	<p>causes Oracle to fire the trigger before executing the triggering event. For row triggers, this is a separate firing before each affected row is changed.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify a <code>BEFORE</code> trigger on a view or an object view.</li> <li>■ When defining a <code>BEFORE</code> trigger for LOB columns, you can read the <code>:OLD</code> value but not the <code>:NEW</code> value. You cannot write either the <code>:OLD</code> or the <code>:NEW</code> value.</li> </ul>
AFTER	<p>causes Oracle to fire the trigger after executing the triggering event. For row triggers, this is a separate firing after each affected row is changed.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot specify an <code>AFTER</code> trigger on a view or an object view.</li> <li>■ When defining an <code>AFTER</code> trigger for LOB columns, you can read the <code>:OLD</code> value but not the <code>:NEW</code> value. You cannot write either the <code>:OLD</code> or the <code>:NEW</code> value.</li> </ul>



---

**Note:** When you create a materialized view log for a table, Oracle implicitly creates an `AFTER ROW` trigger on the table. This trigger inserts a row into the materialized view log whenever an `INSERT`, `UPDATE`, or `DELETE` statement modifies the table's data. You cannot control the order in which multiple row triggers fire. Therefore, you should not write triggers intended to affect the content of the materialized view.

**See Also:** "[CREATE MATERIALIZED VIEW LOG / SNAPSHOT LOG](#)" on page 7-333 for more information on materialized view logs.

---

`INSTEAD OF` causes Oracle to fire the trigger instead of executing the triggering event. By default, `INSTEAD OF` triggers are activated for each row.

If a view is inherently updatable and has `INSTEAD OF` triggers, the triggers take preference. In other words, Oracle fires the triggers instead of performing DML on the view.

**Restrictions:**

- `INSTEAD OF` is a valid clause only for views. You cannot specify an `INSTEAD OF` trigger on a table.
- If a view has `INSTEAD OF` triggers, any views created on it must have `INSTEAD OF` triggers, even if the views are inherently updatable.
- When defining `INSTEAD OF` triggers for LOB columns, you can read both the `:OLD` and the `:NEW` value, but you cannot write either the `:OLD` or the `:NEW` values.

---

**Note:** You can create multiple triggers of the same type (`BEFORE`, `AFTER`, or `INSTEAD OF`) that fire for the same statement on the same table. The order in which Oracle fires these triggers is indeterminate. If your application requires that one trigger be fired before another of the same type for the same statement, combine these triggers into a single trigger whose trigger action performs the trigger actions of the original triggers in the appropriate order.

---

<i>dml_event_clause</i>	specifies one of three DML statements that can cause the trigger to fire. Oracle fires the trigger in the existing user transaction.
<code>DELETE</code>	causes Oracle to fire the trigger whenever a <code>DELETE</code> statement removes a row from the table or an element from a nested table.
<code>INSERT</code>	causes Oracle to fire the trigger whenever an <code>INSERT</code> statement adds a row to table or an element to a nested table.
<code>UPDATE</code>	causes Oracle to fire the trigger whenever an <code>UPDATE</code> statement changes a value in one of the columns specified after <code>OF</code> . If you omit <code>OF</code> , Oracle fires the trigger whenever an <code>UPDATE</code> statement changes a value in any column of the table or nested table.

For an `UPDATE` trigger, you can specify object type, varray, and `REF` columns after `OF` to indicate that the trigger should be fired whenever an `UPDATE` statement changes a value in one of the columns. However, you cannot change the values of these columns in the body of the trigger itself.

---

**Note:** Using OCI functions or the `DBMS_LOB` package to update LOB values or LOB attributes of object columns does not cause Oracle to fire triggers defined on the table containing the columns or the attributes.

---

**Restrictions:**

- You cannot specify `OF` with `UPDATE` for an `INSTEAD OF` trigger. Oracle fires `INSTEAD OF` triggers whenever an `UPDATE` changes a value in any column of the view.
- You cannot specify nested table or LOB columns with `OF`.
- See *AS subquery* of "**CREATE VIEW**" on page 7-456 for a list of constructs that prevent inserts, updates, or deletes on a view.

Performing DML operations directly on nested table columns does not cause Oracle to fire triggers defined on the table containing the nested table column

*ddl\_event*

is one or more types of DDL statements that can cause the trigger to fire. You can create triggers for these events on `DATABASE` or `SCHEMA` unless otherwise noted. You can create `BEFORE` and `AFTER` triggers for these events. Oracle fires the trigger in the existing user transaction. The following values are valid:

<code>ALTER</code>	fires the trigger whenever an <code>ALTER</code> statement modifies a database object in the data dictionary. <b>Restriction:</b> The trigger will not be fired by an <code>ALTER DATABASE</code> statement.
<code>ANALYZE</code>	fires the trigger whenever Oracle collects or deletes statistics or validates the structure of a database object.
<code>ASSOCIATE STATISTICS</code>	fires the trigger whenever Oracle associates a statistics type with a database object.
<code>AUDIT</code>	fires the trigger whenever Oracle tracks the occurrence of a SQL statement or tracks operations on a schema object.
<code>COMMENT</code>	fires the trigger whenever a comment on a database object is added to the data dictionary.
<code>CREATE</code>	fires the trigger whenever a <code>CREATE</code> statement adds a new database object to the data dictionary. <b>Restriction:</b> The trigger will not be fired by a <code>CREATE DATABASE</code> or <code>CREATE CONTROLFILE</code> statement.
<code>DISASSOCIATE STATISTICS</code>	fires the trigger whenever Oracle disassociates a statistics type from a database object.
<code>DROP</code>	fires the trigger whenever a <code>DROP</code> statement removes a database object from the data dictionary.
<code>GRANT</code>	fires the trigger whenever a user grants system privileges or roles or object privileges to another user or to a role.

---

NOAUDIT	fires the trigger whenever a NOAUDIT statement instructs Oracle to stop tracking a SQL statement or operations on a schema object.
RENAME	fires the trigger whenever a RENAME statement change the name of a database object.
REVOKE	fires the trigger whenever a REVOKE statement removes system privileges or roles or object privileges from a user or role.
TRUNCATE	fires the trigger whenever a TRUNCATE statement removes the rows from a table or cluster and resets its storage characteristics.
DDL	fires the trigger whenever any of the preceding DDL statements is issued.
	<b>Restriction:</b> You cannot specify as a triggering event any DDL operation performed through a PL/SQL procedure.
<i>database_event</i>	describes a particular state of the database that can cause the trigger to fire. You can create triggers for these events on DATABASE or SCHEMA unless otherwise noted. For each of these triggering events, Oracle opens an autonomous transaction scope, fires the trigger, and commits any separate transaction (regardless of any existing user transaction). <b>See Also:</b> <i>PL/SQL User's Guide and Reference</i> for more information on autonomous transaction scope.
SERVERERROR	causes Oracle to fire the trigger whenever a server error message is logged. <b>Note:</b> The following errors do not cause a SERVERERROR trigger to fire: <ul style="list-style-type: none"> <li>■ ORA-01403: data not found</li> <li>■ ORA-01422: exact fetch returns more than requested number of rows</li> <li>■ ORA-01423: error encountered while checking for extra rows in exact fetch</li> <li>■ ORA-01034: ORACLE not available</li> <li>■ ORA-04030: out of process memory</li> </ul>
LOGON	causes Oracle to fire the trigger whenever a client application logs onto the database.
LOGOFF	causes Oracle to fire the trigger whenever a client applications logs off the database.
STARTUP	causes Oracle to fire the trigger whenever the database is opened.
SHUTDOWN	causes Oracle to fire the trigger whenever an instance of the database is shut down.

**Notes:**

- Only AFTER triggers are relevant for LOGON, STARTUP, and SERVERERROR.
- Only BEFORE triggers are relevant for LOGOFF and SHUTDOWN.
- AFTER STARTUP and BEFORE SHUTDOWN triggers apply only to DATABASE.

ON	determines the database object on which the trigger is to be created.
<i>[schema.] table / view</i>	<p>specifies the <i>schema</i> and <i>table</i> or <i>view</i> name of one of the following on which the trigger is to be created:</p> <ul style="list-style-type: none"> <li>■ table or view</li> <li>■ object table or object view</li> <li>■ a column of nested-table type</li> </ul> <p>If you omit <i>schema</i>, Oracle assumes the table is in your own schema. You can create triggers on index-organized tables.</p> <p><b>Restriction:</b> You cannot create a trigger on a table in the schema SYS.</p>
NESTED TABLE	<p>specifies that the trigger is being defined on column <i>nested_table_column</i> of a view. Such a trigger will fire only if the DML operates on the elements of the nested table.</p> <p><b>Restriction:</b> You can specify NESTED TABLE only for INSTEAD OF triggers.</p>
DATABASE	specifies that the trigger is being defined on the entire database.
SCHEMA	specifies that the trigger is being defined on the current schema.
<i>referencing_ clause</i>	<p>specifies correlation names. You can use correlation names in the PL/SQL block and WHEN condition of a row trigger to refer specifically to old and new values of the current row. The default correlation names are OLD and NEW. If your row trigger is associated with a table named OLD or NEW, use this clause to specify different correlation names to avoid confusion between the table name and the correlation name.</p> <ul style="list-style-type: none"> <li>■ If the trigger is defined on a <b>nested table</b>, OLD and NEW refer to the row of the nested table, and PARENT refers to the current row of the parent table.</li> <li>■ If the trigger is defined on an object table or view, OLD and NEW refer to object instances.</li> </ul> <p><b>Restriction:</b> This clause is valid only for DML event triggers (not DDL or database event triggers).</p>
FOR EACH ROW	<p>designates the trigger to be a row trigger. Oracle fires a row trigger once for each row that is affected by the triggering statement and meets the optional trigger constraint defined in the WHEN condition.</p> <p><b>Note:</b> This clause is applies only to DML events, not to DDL or database events.</p>

---

Except for `INSTEAD OF` triggers, if you omit this clause, the trigger is a statement trigger. Oracle fires a statement trigger only once when the triggering statement is issued if the optional trigger constraint is met.

`INSTEAD OF` trigger statements are implicitly activated for each row.

`WHEN (condition)` specifies the trigger restriction, which is a SQL condition that must be satisfied for Oracle to fire the trigger. See the syntax description of condition in "[Conditions](#)" on page 5-14. This condition must contain correlation names and cannot contain a query.

**Restrictions:**

- You can specify a trigger restriction only for a row trigger. Oracle evaluates this condition for each row affected by the triggering statement.
- You cannot specify trigger restrictions for `INSTEAD OF` trigger statements.
- You can reference object columns or their attributes, or varray, nested table, or LOB columns. You cannot invoke PL/SQL functions or methods in the trigger restriction.

`pl/sql_block` is the PL/SQL block that Oracle executes to fire the trigger. For information on PL/SQL, including how to write PL/SQL blocks, see *PL/SQL User's Guide and Reference*.

The PL/SQL block of a database trigger can contain one of a series of built-in functions in the SYS schema designed solely to extract system event attributes. These functions can be used **only** in the PL/SQL block of a database trigger. For information on these functions, see *Oracle8i Application Developer's Guide - Fundamentals*.

**Restrictions:**

- The PL/SQL block of a trigger cannot contain transaction control SQL statements (`COMMIT`, `ROLLBACK`, `SAVEPOINT`, and `SET CONSTRAINT`) if the block is executed within the same transaction.
- You can reference and use LOB columns in the trigger action inside the PL/SQL block, but you cannot modify their values within the trigger action.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals*.

`call_procedure_statement` enables you to call a stored procedure, rather than specifying inline the trigger code as a PL/SQL block. The syntax of this statement is the same as that for "[CALL](#)" on page 7-226, with the following exceptions:

- You cannot specify the `INTO` clause of `CALL`, because it applies only to functions.
  - You cannot specify bind variables in `expr`.
  - To reference columns of tables on which the trigger is being defined, you must specify `:NEW` and `:OLD`. See the "[Calling a Procedure in a Trigger Body Example](#)" on page 7-434.
-

## Examples

**DML Trigger Example** This example creates a `BEFORE` statement trigger named `EMP_PERMIT_CHANGES` in the schema `SCOTT`. You would write such a trigger to place restrictions on DML statements issued on this table (such as when such statements could be issued).

```
CREATE TRIGGER scott.emp_permit_changes
  BEFORE
  DELETE OR INSERT OR UPDATE
  ON scott.emp
  pl/sql block
```

Oracle fires this trigger whenever a `DELETE`, `INSERT`, or `UPDATE` statement affects the `EMP` table in the schema `SCOTT`. The trigger `EMP_PERMIT_CHANGES` is a `BEFORE` statement trigger, so Oracle fires it once before executing the triggering statement.

**DML Trigger Example with Restriction** This example creates a `BEFORE` row trigger named `SALARY_CHECK` in the schema `SCOTT`. The PL/SQL block might specify, for example, that the employee's salary must fall within the established salary range for the employee's job:

```
CREATE TRIGGER scott.salary_check
  BEFORE
  INSERT OR UPDATE OF sal, job ON scott.emp
  FOR EACH ROW
  WHEN (new.job <> 'PRESIDENT')
  pl/sql_block
```

Oracle fires this trigger whenever one of the following statements is issued:

- an `INSERT` statement that adds rows to the `EMP` table
- an `UPDATE` statement that changes values of the `SAL` or `JOB` columns of the `EMP` table

`SALARY_CHECK` is a `BEFORE` row trigger, so Oracle fires it before changing each row that is updated by the `UPDATE` statement or before adding each row that is inserted by the `INSERT` statement.

`SALARY_CHECK` has a trigger restriction that prevents it from checking the salary of the company president.

**Calling a Procedure in a Trigger Body Example** You could create the `SALARY_CHECK` trigger described in the preceding example by calling a procedure instead of

providing the trigger body in a PL/SQL block. Assume you have defined a procedure SCOTT.CHECK\_SAL, which verifies that an employee's salary is in an appropriate range. Then you could create the trigger SALARY\_CHECK as follows:

```
CREATE TRIGGER scott.salary_check
  BEFORE INSERT OR UPDATE OF sal, job ON scott.emp
  FOR EACH ROW
  WHEN (new.job <> 'PRESIDENT')
  CALL check_sal(:new.job, :new.sal, :new.ename)
```

The procedure CHECK\_SAL could be implemented in PL/SQL, C, or Java. Also, you can specify :OLD values in the CALL clause instead of :NEW values.

**Database Event Trigger Example** This example creates a trigger to log all errors. The PL/SQL block does some special processing for a particular error (invalid logon, error number 1017). This trigger is an AFTER statement trigger, so it is fired after an unsuccessful statement execution (such as unsuccessful logon).

```
CREATE TRIGGER log_errors AFTER SERVERERROR ON DATABASE
  BEGIN
    IF (IS_SERVERERROR (1017)) THEN
      <special processing of logon error>
    ELSE
      <log error number>
    END IF;
  END;
```

**DDL Trigger Example** This example creates an AFTER statement trigger on any DDL statement CREATE. Such a trigger can be used to audit the creation of new data dictionary objects in your schema.

```
CREATE TRIGGER audit_db_object AFTER CREATE
  ON SCHEMA
  pl/sql_block
```

**INSTEAD OF Trigger Example** In this example, customer data is stored in two tables. The object view ALL\_CUSTOMERS is created as a UNION of the two tables, CUSTOMERS\_SJ and CUSTOMERS\_PA. An INSTEAD OF trigger is used to insert values.

```
CREATE TABLE customers_sj
  ( cust      NUMBER(6),
    address  VARCHAR2(50),
    credit   NUMBER(9,2) );
```

## CREATE TRIGGER

---

```
CREATE TABLE customers_pa
( cust      NUMBER(6),
  address  VARCHAR2(50),
  credit   NUMBER(9,2) );

CREATE TYPE customer_t AS OBJECT
( cust      NUMBER(6),
  address   VARCHAR2(50),
  credit    NUMBER(9,2),
  location  VARCHAR2(20) );

CREATE VIEW all_customers (cust)
AS SELECT customer_t (cust, address, credit, 'SAN_JOSE')
FROM   customers_sj
UNION ALL
SELECT customer_t (cust, address, credit, 'PALO_ALTO')
FROM   customers_pa;


CREATE TRIGGER instrig INSTEAD OF INSERT ON all_customers
FOR EACH ROW
BEGIN
  IF (:new.cust.location = 'SAN_JOSE') THEN
    INSERT INTO customers_sj
      VALUES (:new.cust.cust, :new.cust.address, :new.cust.credit);
  ELSE
    INSERT INTO customers_pa
      VALUES (:new.cust.cust, :new.cust.address, :new.cust.credit);
  END IF;
END;
```




## CREATE TYPE

### Syntax

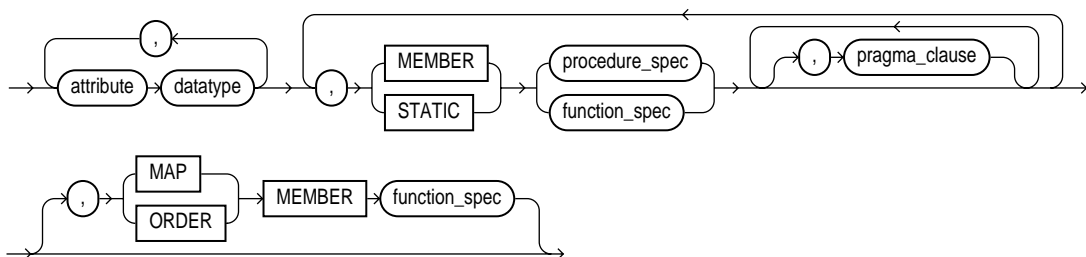
**create\_incomplete\_type::=**




**create\_object\_type::=**



**element\_list::=**

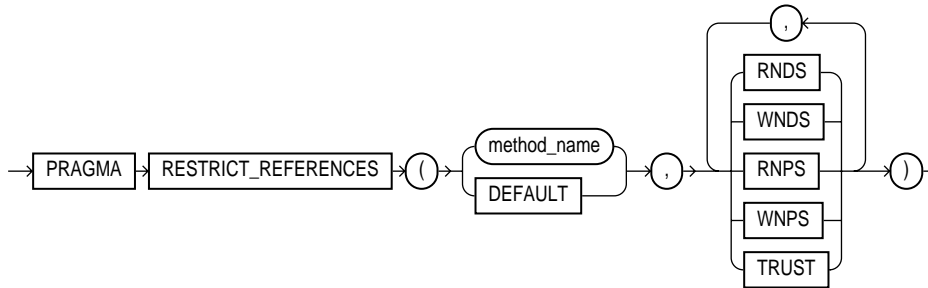


**invoker\_rights\_clause::=**




## CREATE TYPE


### pragma\_clause::=




### procedure\_spec | function\_spec::=




### call\_spec::=





### Java\_declaration::=



### C\_declaration::=



**create\_varray\_type::=****create\_nested\_table\_type::=****Purpose**

To create an **object type**, named varying array (**varray**), **nested table type**, or an **incomplete object type**.

Oracle implicitly defines a constructor method for each user-defined type that you create. A **constructor** is a system-supplied procedure that is used in SQL statements or in PL/SQL code to construct an instance of the type value. The name of the constructor method is the same as the name of the user-defined type.

The parameters of the object type constructor method are the data attributes of the object type. They occur in the same order as the attribute definition order for the object type. The parameters of a nested table or varray constructor are the elements of the nested table or the varray.

An **incomplete type** is a type created by a forward type definition. It is called "incomplete" because it has a name but no attributes or methods. It can be referenced by other types, and so can be used to define types that refer to each other. However, you must fully specify the type before you can use it to create a table or an object column or a column of a nested table type.

**See Also:** *PL/SQL User's Guide and Reference, Oracle8i Application Developer's Guide - Fundamentals, and Oracle8i Concepts* for more information about objects, incomplete types, varrays, and nested tables.

## Prerequisites

To create a type in your own schema, you must have the `CREATE TYPE` system privilege. To create a type in another user's schema, you must have the `CREATE ANY TYPE` system privilege. You can acquire these privileges explicitly or be granted them through a role.

The owner of the type must either be explicitly granted the `EXECUTE` object privilege in order to access all other types referenced within the definition of the type, or the type owner must be granted the `EXECUTE ANY TYPE` system privilege. The owner **cannot** obtain these privileges through roles.

If the type owner intends to grant other users access to the type, the owner must be granted the `EXECUTE` object privilege to the referenced types with the `GRANT OPTION` or the `EXECUTE ANY TYPE` system privilege with the `ADMIN OPTION`. Otherwise, the type owner has insufficient privileges to grant access on the type to other users.

## Keywords and Parameters

---

<code>OR REPLACE</code>	re-creates the type if it already exists. Use this clause to change the definition of an existing type without first dropping it.  Users previously granted privileges on the re-created object type can use and reference the object type without being granted privileges again.  If any function-based indexes depend on the type, Oracle marks the indexes <code>DISABLED</code> .
<code>schema</code>	is the schema to contain the type. If you omit <code>schema</code> , Oracle creates the type in your current schema.
<code>type_name</code>	is the name of an object type, a nested table type, or a varray type.  If creating the type results in compilation errors, Oracle returns an error. You can see the associated compiler error messages with the SQL*Plus command <code>SHOW ERRORS</code> .
<code>create_object_type</code>	creates the type as a user-defined object type. The variables that form the data structure are called <i>attributes</i> . The member subprograms that define the object's behavior are called <i>methods</i> . <code>AS OBJECT</code> is required when creating an object type.

*invoker\_rights\_clause* specifies whether the member functions and procedures of the object type execute with the privileges and in the schema of the user who owns the object type or with the privileges and in the schema of `CURRENT_USER`. This specification applies to the corresponding type body as well. (For information on how `CURRENT_USER` is determined, see *Oracle8i Concepts and Oracle8i Application Developer's Guide - Fundamentals*.)

This clause also determines how Oracle resolves external names in queries, DML operations, and dynamic SQL statements in the member functions and procedures of the type.

**Restriction:** You can specify this clause only for an object type, not for a nested table or varray type.

**See Also:** *PL/SQL User's Guide and Reference*.

`AUTHID CURRENT_USER` specifies that the member functions and procedures of the object type execute with the privileges of `CURRENT_USER`. This clause creates an "invoker-rights type."

This clause also specifies that external names in queries, DML operations, and dynamic SQL statements resolve in the schema of `CURRENT_USER`. External names in all other statements resolve in the schema in which the type resides.

`AUTHID DEFINER` specifies that the member functions and procedures of the object type execute with the privileges of the owner of the schema in which the functions and procedures reside, and that external names resolve in the schema where the member functions and procedures reside. This is the default.

*datatype* is the name of the attribute's Oracle built-in datatype or user-defined type. For a list of possible datatypes, see "[Datatypes](#)" on page 2-9.

**Restrictions:**

- You cannot specify attributes of type `ROWID`, `LONG`, or `LONG ROW`.
- You cannot create an object with `NCLOB`, `NCHAR`, or `NVARCHAR2` attributes, but you can specify parameters of these datatypes in methods.
- You cannot specify a datatype of `UROWID` for a user-defined object type.
- If you specify an object of type `REF`, the target object must have an object identifier.

*attribute* specifies, for an object type, the name of an object attribute. Attributes are data items with a name and a type specifier that form the structure of the object. You must specify at least one attribute for each object type.

`MEMBER` specifies a function or procedure subprogram associated with the object type that is referenced as an attribute. Typically, you invoke member methods in a "selfish" style, such as `object_expression.method()`. This class of method has an implicit first argument referenced as `SELF` in the method's body, which represents the object on which the method has been invoked.

**STATIC** also specifies a function or procedure subprogram associated with the object type. However, unlike member methods, static methods do not have any implicit parameters (that is, `SELF` is not referenceable in their body). They are typically invoked as `type_name.method()`.

For both member and static methods, you must specify a corresponding method body in the object type body for each procedure or function specification. See ["CREATE TYPE BODY"](#) on page 7-447. For information about method invocation and methods, see *PL/SQL User's Guide and Reference*.

*procedure\_spec* | *function\_spec* is the specification of a procedure or function subprogram. The `RETURN` clause is valid only for a function. The syntax shown is an abbreviated form. For the full syntax with all possible clauses, see ["CREATE PROCEDURE"](#) on page 7-353 and ["CREATE FUNCTION"](#) on page 7-284.

If this subprogram does not include the declaration of the procedure or function, you must issue a corresponding `CREATE TYPE BODY` statement. See ["CREATE TYPE BODY"](#) on page 7-447.

For a list of restrictions on user-defined functions, see ["Restrictions on User-Defined Functions"](#) on page 7-286.

*call\_spec* is the call specification ("call spec") that maps a Java or C method name, parameter types, and return type to their SQL counterparts. If all the member methods in the type have been defined in this clause, you need not issue a corresponding `CREATE TYPE BODY` statement.

In *Java\_declaration*, *'string'* identifies the Java implementation of the method.

**See Also:**

- *Oracle8i Java Stored Procedures Developer's Guide*.
- *Oracle8i Application Developer's Guide - Fundamentals* for an explanation of the parameters and semantics of the *C\_declaration*.

*pragma\_clause* specifies a compiler directive.

**PRAGMA RESTRICT\_ REFERENCES** is a compiler directive that denies member functions read/write access to database tables, packaged variables, or both, and thereby helps to avoid side effects.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals*.

*method\_name* is the name of the `MEMBER` function or procedure to which the pragma is being applied.

**DEFAULT** specifies that the pragma should be applied to all methods in the type for which a pragma has not been explicitly specified.

**WNDS** specifies the constraint *writes no database state* (does not modify database tables).

**WNPS** specifies the constraint *writes no package state* (does not modify packaged variables).

**RNDS** specifies the constraint *reads no database state* (does not query database tables).

---

	RNPS	specifies the constraint <i>reads no package state</i> (does not reference package variables).
	TRUST	specifies that the restrictions listed in the pragma are not actually to be enforced, but are simply trusted to be true.
MAP MEMBER <i>function_spec</i>		<p>specifies a member function (map method) that returns the relative position of a given instance in the ordering of all instances of the object. A map method is called implicitly and induces an ordering of object instances by mapping them to values of a predefined <i>scalar</i> type. PL/SQL uses the ordering to evaluate Boolean expressions and to perform comparisons.</p> <p>If the argument to the map method is null, the map method returns null and the method is not invoked.</p> <p>An object specification can contain only one map method, which must be a function. The result type must be a predefined SQL scalar type, and the map function can have no arguments other than the implicit <i>SELF</i> argument.</p> <hr/> <p><b>Note:</b> If <i>type_name</i> will be referenced in queries involving sorts (through an <i>ORDER BY</i>, <i>GROUP BY</i>, <i>DISTINCT</i>, or <i>UNION</i> clause) or joins, and you want those queries to be parallelized, you must specify a <i>MAP</i> member function.</p>
ORDER MEMBER <i>function_spec</i>		<p>specifies a member function (<i>ORDER</i> method) that takes an instance of an object as an explicit argument and the implicit <i>SELF</i> argument and returns either a negative, zero, or positive integer. The negative, positive, or zero indicates that the implicit <i>SELF</i> argument is less than, equal to, or greater than the explicit argument.</p> <p>If either argument to the order method is null, the order method returns null and the method is not invoked.</p> <p>When instances of the same object type definition are compared in an <i>ORDER BY</i> clause, the order method <i>function_specification</i> is invoked.</p> <p>An object specification can contain only one <i>ORDER</i> method, which must be a function having the return type <i>NUMBER</i>.</p>

You can define either a *MAP* method or an *ORDER* method in a type specification, but not both. If you declare either method, you can compare object instances in SQL.

If neither a *MAP* nor an *ORDER* method is specified, only comparisons for equality or inequality can be performed. Therefore object instances cannot be ordered. Instances of the same type definition are equal only if each pair of their corresponding attributes is equal. No comparison method needs to be specified to determine the equality of two object types.

Use *MAP* if you are performing extensive sorting or hash join operations on object instances. *MAP* is applied once to map the objects to scalar values and then the scalars are used during sorting and merging. A *MAP* method is more efficient than an *ORDER* method, which must invoke the method for each object comparison. You must use a *MAP* method for hash joins. You cannot use an *ORDER* method because the hash mechanism hashes on the object value.

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals* for more information about object value comparisons.

*create\_varray\_type* creates the type as an ordered set of elements, each of which has the same datatype. You must specify a name and a maximum *limit* of zero or more. The array limit must be an integer literal. Oracle does not support anonymous varrays.

The type name for the objects contained in the varray must be one of the following:

- A built-in datatype,
- A REF, or
- An object type.

**Restrictions:**

- A collection type cannot contain any other collection type, either directly or indirectly. That is, a varray type cannot contain any elements that are or contain varrays or nested tables.
- You cannot create varray types of LOB datatypes.

*create\_nested\_table\_type* creates a named nested table of type *datatype*.

When *datatype* is an object type, the nested table type describes a table whose columns match the name and attributes of the object type.

When *datatype* is a scalar type, then the nested table type describes a table with a single, scalar type column called "column\_value".

**Restrictions:**

- A collection type cannot contain any other collection type, either directly or indirectly. That is, a nested table type cannot contain any elements that are or contain varrays or nested tables.
  - You cannot specify NCLOB for *datatype*. However, you can specify CLOB or BLOB.
- 

## Examples

**Object Type Example** The following example creates object type PERSON\_T with LOB attributes:

```
CREATE TYPE person_t AS OBJECT
  (name    CHAR(20),
   resume  CLOB,
   picture BLOB);
```

**Varray Type Example** The following statement creates MEMBERS\_TYPE as a varray type with 100 elements:

```
CREATE TYPE members_type AS VARRAY(100) OF CHAR(5);
```



**Nested Table Type Example** The following example creates a named table type PROJECT\_TABLE of object type PROJECT\_T:

```
CREATE TYPE project_t AS OBJECT
  (pno CHAR(5),
   pname CHAR(20),
   budgets DEC(7,2));

CREATE TYPE project_table AS TABLE OF project_t;
```

**Constructor Example** The following example invokes method constructor COL.GETBAR():

```
CREATE TYPE foo AS OBJECT (a1 NUMBER,
                          MEMBER FUNCTION getbar RETURN NUMBER,);
CREATE TABLE footab(col foo);

SELECT col.getbar() FROM footab;
```

Unlike function invocations, method invocations require parentheses, even when the methods do not have additional arguments.

The next example invokes the system-defined constructor to construct the FOO\_T object and insert it into the FOO\_TAB table:

```
CREATE TYPE foo_t AS OBJECT (a1 NUMBER, a2 NUMBER);
CREATE TABLE foo_tab (b1 NUMBER, b2 foo_t);
INSERT INTO foo_tab VALUES (1, foo_t(2,3));
```

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals and PL/SQL User's Guide and Reference* for more information about constructors.

**Static Method Example** The following example changes the definition of the EMPLOYEE\_T type to associate it with the CONSTRUCT\_EMP function:

```
CREATE OR REPLACE TYPE employee_t AS OBJECT(
  empid RAW(16),
  ename CHAR(31),
  dept REF department_t,
  STATIC function construct_emp
    (name VARCHAR2, dept REF department_t)
  RETURN employee_t
);
```

This statement requires the following type body statement:

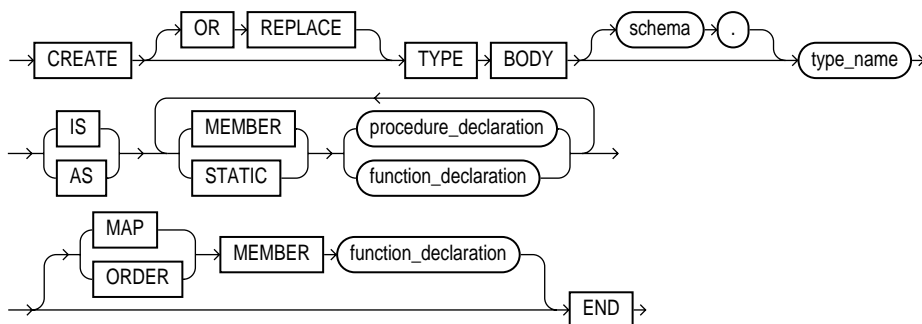
```
CREATE OR REPLACE TYPE BODY employee_t IS
  STATIC FUNCTION construct_emp
    (name varchar2, dept REF department_t)
  RETURN employee_t IS
  BEGIN
    return employee_t(SYS_GUID(),name,dept);
  END;
END;
```

This type and type body definition allows the following operation:


```
INSERT INTO emptab
  VALUES (employee_t.construct_emp('John Smith', NULL));
```

## CREATE TYPE BODY


### Syntax




### `procedure_declaration | function_declaration ::=`




### `call_spec ::=`



### `Java_declaration ::=`



### `C_declaration ::=`



## Purpose

To define or implement the member methods defined in the object type specification. You create object types with the `CREATE TYPE` and the `CREATE TYPE BODY` statements. The `CREATE TYPE` statement specifies the name of the object type, its attributes, methods, and other properties. The `CREATE TYPE BODY` statement contains the code for the methods in the type.

For each method specified in an object type specification for which you did not specify the *call\_spec*, you must specify a corresponding method body in the object type body.

For information on creating and modifying a type specification, see "[CREATE TYPE](#)" on page 7-437 and "[ALTER TYPE](#)" on page 7-186.

## Prerequisites

Every member declaration in the `CREATE TYPE` specification for object types must have a corresponding construct in the `CREATE TYPE` or `CREATE TYPE BODY` statement.

To create or replace a type body in your own schema, you must have the `CREATE TYPE` or the `CREATE ANY TYPE` system privilege. To create an object type in another user's schema, you must have the `CREATE ANY TYPE` system privileges. To replace an object type in another user's schema, you must have the `DROP ANY TYPE` system privileges.

## Keywords and Parameters

---

<code>OR REPLACE</code>	<p>re-creates the type body if it already exists. Use this clause to change the definition of an existing type body without first dropping it.</p> <p>Users previously granted privileges on the re-created object type body can use and reference the object type body without being granted privileges again.</p> <p>You can use this clause to add new member subprogram definitions to specifications added with the <code>ALTER TYPE . . . REPLACE</code> statement.</p>
<i>schema</i>	is the schema to contain the type body. If you omit <i>schema</i> , Oracle creates the type body in your current schema.
<i>type_name</i>	is the name of an object type.
<code>MEMBER</code>   <code>STATIC</code>	declares or implements a method function or procedure subprogram associated with the object type specification. For a description of the difference between member and static methods, see " <a href="#">CREATE TYPE</a> " on page 7-437. For information about overloading subprogram names within a package, see <i>PL/SQL User's Guide and Reference</i> .

You must define a corresponding method name, optional parameter list, and (for functions) a return type in the object type specification for each procedure or function declaration.

*procedure\_ declaration* is the declaration of a procedure subprogram.

*function\_ declaration* is the declaration of a function subprogram.

**See Also:** "[CREATE PROCEDURE](#)" on page 7-353, "[CREATE FUNCTION](#)" on page 7-284, and *Oracle8i Application Developer's Guide - Fundamentals*.

**MAP MEMBER** declares or implements a member function (MAP method) that returns the relative position of a given instance in the ordering of all instances of the object. A map method is called implicitly and specifies an ordering of object instances by mapping them to values of a predefined scalar type. PL/SQL uses the ordering to evaluate Boolean expressions and to perform comparisons.

If the argument to the map method is null, the map method returns null and the method is not invoked.

An object type body can contain only one map method, which must be a function. The map function can have no arguments other than the implicit SELF argument.

**ORDER MEMBER** specifies a member function (ORDER method) that takes an instance of an object as an explicit argument and the implicit SELF argument and returns either a negative, zero, or positive integer. The negative, positive, or zero indicates that the implicit SELF argument is less than, equal to, or greater than the explicit argument.

If either argument to the order method is null, the order method returns null and the method is not invoked.

When instances of the same object type definition are compared in an ORDER BY clause, Oracle invokes the order method *function\_spec*.

An object specification can contain only one ORDER method, which must be a function having the return type NUMBER.

You can declare either a MAP method or an ORDER method, but not both. If you declare either method, you can compare object instances in SQL.

If you do not declare either method, you can compare object instances only for equality or inequality. Instances of the same type definition are equal only if each pair of their corresponding attributes is equal.

*procedure\_ declaration* | *function\_ declaration* is the declaration of a procedure or function subprogram. The RETURN clause is valid only for a function. The syntax shown is an abbreviated form. For the full syntax with all possible clauses, see "[CREATE PROCEDURE](#)" on page 7-353 and "[CREATE FUNCTION](#)" on page 7-284.

*pl/sql\_block* declares the procedure or function.

**See Also:** *PL/SQL User's Guide and Reference*.

---

<i>call_spec</i>	<p>is the call specification ("call spec") that maps a Java or C method name, parameter types, and return type to their SQL counterparts. In <i>Java_declaration</i>, 'string' identifies the Java implementation of the method.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"><li>▪ <i>Oracle8i Java Stored Procedures Developer's Guide</i>.</li><li>▪ <i>Oracle8i Application Developer's Guide - Fundamentals</i> for an explanation of the parameters and semantics of the <i>C_declaration</i>.</li></ul>
AS EXTERNAL	<p>is an alternative way of declaring a C method. This clause has been deprecated and is supported for backward compatibility only. Oracle Corporation recommends that you use the <i>call_spec</i> syntax with the <i>C_declaration</i>.</p>

---

## Examples

The following object type body implements member subprograms for RATIONAL.

```
CREATE TYPE BODY rational
IS
  MAP MEMBER FUNCTION rat_to_real RETURN REAL IS
    BEGIN
      RETURN numerator/denominator;
    END;

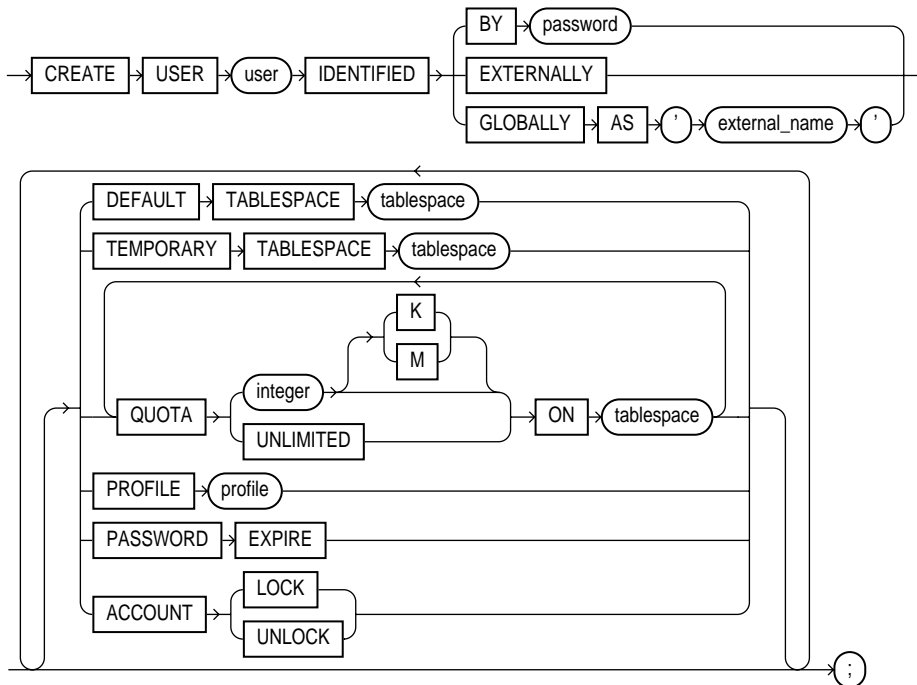
  MEMBER PROCEDURE normalize IS
    gcd NUMBER := integer_operations.greatest_common_divisor
      (numerator, denominator);
    BEGIN
      numerator := numerator/gcd;
      denominator := denominator/gcd;
    END;

  MEMBER FUNCTION plus(x rational) RETURN rational IS
    r rational := rational_operations.make_rational
      (numerator*x.denominator +
       x.numerator*denominator,
       denominator*x.denominator);
    BEGIN
      RETURN r;
    END;

END;
```

## CREATE USER

### Syntax



### Purpose

To create and configure a database **user**, or an account through which you can log in to the database and establish the means by which Oracle permits access by the user.

---

**Note:** You can enable a user to connect to Oracle through a proxy (that is, an application or application server). For syntax and discussion, refer to "[ALTER USER](#)" on page 7-193.

---

## Prerequisites

You must have `CREATE USER` system privilege. When you create a user with the `CREATE USER` statement, the user's privilege domain is empty. To log on to Oracle, a user must have `CREATE SESSION` system privilege. Therefore, after creating a user, you should grant the user at least the `CREATE SESSION` privilege. See "[GRANT system\\_privileges\\_and\\_roles](#)" on page 7-519.

## Keywords and Parameters

---

<code>user</code>	is the name of the user to be created. This name can contain only characters from your database character set and must follow the rules described in the section " <a href="#">Schema Object Naming Rules</a> " on page 2-72. Oracle recommends that the user name contain at least one single-byte character regardless of whether the database character set also contains multi-byte characters.
<code>IDENTIFIED</code>	indicates how Oracle authenticates the user.  <b>See Also:</b> <i>Oracle8i Application Developer's Guide - Fundamentals</i> and your operating system specific documentation for more information.
<code>BY password</code>	creates a <b>local user</b> and indicates that the user must specify <i>password</i> to log on. Passwords can contain only single-byte characters from your database character set regardless of whether this character set also contains multibyte characters.  Passwords must follow the rules described in the section " <a href="#">Schema Object Naming Rules</a> " on page 2-72, unless you are using Oracle's password complexity verification routine. That routine requires a more complex combination of characters than the normal naming rules permit. You implement this routine with the <code>UTLPWDMG.SQL</code> script, which is further described in <i>Oracle8i Administrator's Guide</i> .  <b>See Also:</b> <i>Oracle8i Administrator's Guide</i> to for a detailed description and explanation of how to use password management and protection.
<code>EXTERNALLY</code>	creates an <b>external user</b> and indicates that a user must be authenticated by an external service (such as an operating system or a third-party service). Doing so causes Oracle to rely on the login authentication of the operating system to ensure that a specific operating system user has access to a specific database user.  <b>WARNING:</b> Oracle strongly recommends that you do not use <code>IDENTIFIED EXTERNALLY</code> with operating systems that have inherently weak login security. For more information, see <i>Oracle8i Administrator's Guide</i> .
<code>GLOBALLY AS 'external_name'</code>	creates a <b>global user</b> and indicates that a user must be authenticated by the enterprise directory service. The <code>'external_name'</code> string can take one of two forms:



- The X.509 name at the enterprise directory service that identifies this user. It should be of the form 'CN=*username*,*other\_attributes*', where *other\_attributes* is the rest of the user's distinguished name (DN) in the directory.
- A null string ( ' ') indicating that the enterprise directory service will map authenticated global users to the appropriate database schema with the appropriate roles.

**See Also:** *Oracle Advanced Security Administrator's Guide* for more information on global users.

---

**Note:** You can control the ability of an application server to connect as the specified user and to activate that user's roles using the ALTER USER statement.

**See Also:** "ALTER USER" on page 7-193

---

DEFAULT  
TABLESPACE

identifies the default tablespace for objects that the user creates. If you omit this clause, objects default to the SYSTEM tablespace.

**See Also:** "CREATE TABLESPACE" on page 4-419 for more information on tablespaces.

TEMPORARY  
TABLESPACE

identifies the tablespace for the user's temporary segments. If you omit this clause, temporary segments default to the SYSTEM tablespace.

QUOTA

allows the user to allocate space in the tablespace and optionally establishes a quota of *integer* bytes. Use K or M to specify the quota in kilobytes or megabytes. This quota is the maximum space in the tablespace the user can allocate.

A CREATE USER statement can have multiple QUOTA clauses for multiple tablespaces.

UNLIMITED allows the user to allocate space in the tablespace without bound.

PROFILE

reassigns the profile named to the user. The profile limits the amount of database resources the user can use. If you omit this clause, Oracle assigns the DEFAULT profile to the user.

**See Also:** "GRANT system\_privileges\_and\_roles" on page 7-519 and "CREATE PROFILE" on page 7-359.

PASSWORD  
EXPIRE

causes the user's *password* to expire. This setting forces the user (or the DBA) to change the password before the user can log in to the database.

ACCOUNT LOCK

locks the user's account and disables access.

ACCOUNT  
UNLOCK

unlocks the user's account and enables access to the account.

---

## Examples

If you create a new user with PASSWORD EXPIRE, the user's password must be changed before attempting to log in to the database. You can create the user SIDNEY by issuing the following statement:

```
CREATE USER sidney
  IDENTIFIED BY welcome
  DEFAULT TABLESPACE cases_ts
  QUOTA 10M ON cases_ts
  TEMPORARY TABLESPACE temp_ts
  QUOTA 5M ON system
  PROFILE engineer
  PASSWORD EXPIRE;
```

The user SIDNEY has the following characteristics:

- The password WELCOME
- Default tablespace CASES\_TS, with a quota of 10 megabytes
- Temporary tablespace TEMP\_TS
- Access to the tablespace SYSTEM, with a quota of 5 megabytes
- Limits on database resources defined by the profile ENGINEER
- An expired password, which must be changed before SIDNEY can log in to the database

To create a user accessible only by the operating system account GEORGE, prefix GEORGE by the value of the initialization parameter OS\_AUTHENT\_PREFIX. For example, if this value is "OPS\$", you can create the user OPS\$GEORGE with the following statement:

```
CREATE USER ops$george
  IDENTIFIED EXTERNALLY
  DEFAULT TABLESPACE accs_ts
  TEMPORARY TABLESPACE temp_ts
  QUOTA UNLIMITED ON accs_ts;
```

The user OPS\$GEORGE has the following additional characteristics:

- Default tablespace ACCS\_TS
- Default temporary tablespace TEMP\_TS
- Unlimited space on the tablespaces ACCS\_TS and TEMP\_TS
- Limits on database resources defined by the DEFAULT profile

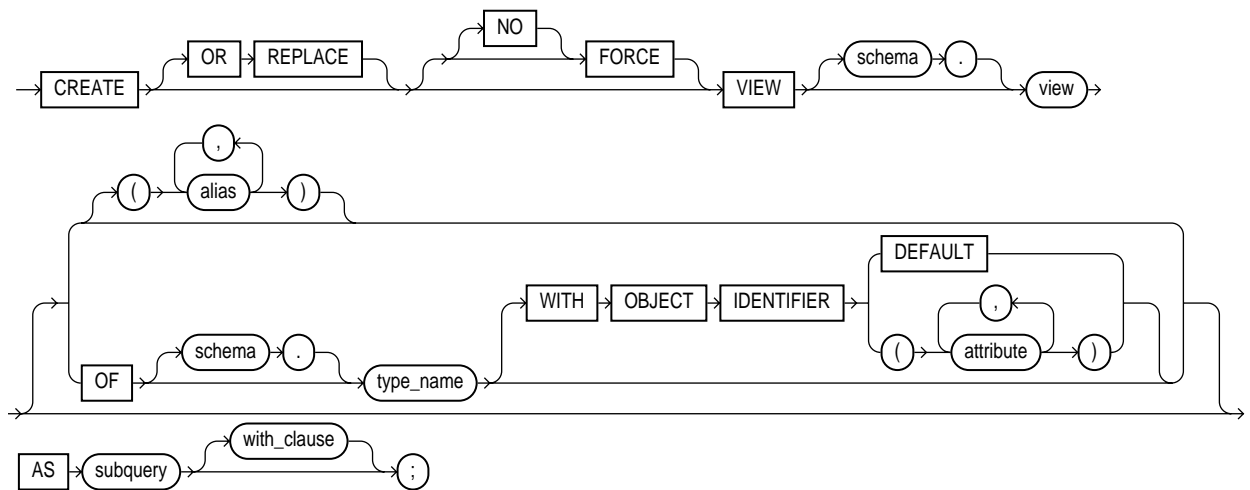
The following example creates user CINDY as a global user:

```
CREATE USER cindy
  IDENTIFIED GLOBALLY AS 'CN=cindy,OU=division1,O=oracle,C=US'
  DEFAULT TABLESPACE legal_ts
```

```
QUOTA 20M ON legal_ts  
PROFILE lawyer;
```


## CREATE VIEW

### Syntax



**subquery:** See ["SELECT and Subqueries"](#) on page 7-569.

**with\_clause::=**



### Purpose

To define a **view**, a logical table based on one or more tables or views. A view contains no data itself. The tables upon which a view is based are called **base tables**.

You can also create an **object view** or a relational view that supports LOB and object datatypes (object types, REFs, nested table, or varray types) on top of the existing view mechanism. An object view is a view of a user-defined type, where each row contains objects, each object with a unique object identifier.

For information on various types of views and their uses, see *Oracle8i Concepts*, *Oracle8i Application Developer's Guide - Fundamentals*, and *Oracle8i Administrator's Guide*.

For information on modifying a view, see "[ALTER VIEW](#)" on page 7-198. For information on removing a view from the database, see "[DROP VIEW](#)" on page 7-511.

## Prerequisites

To create a view in your own schema, you must have `CREATE VIEW` system privilege. To create a view in another user's schema, you must have `CREATE ANY VIEW` system privilege.

The owner of the schema containing the view must have the privileges necessary to either select, insert, update, or delete rows from all the tables or views on which the view is based. For information on these privileges, see "[SELECT and Subqueries](#)" on page 7-569, "[INSERT](#)" on page 7-539, "[UPDATE](#)" on page 7-615, and "[DELETE](#)" on page 7-464. The owner must be granted these privileges directly, rather than through a role.

To use the basic constructor method of an object type when creating an object view, one of the following must be true:

- The object type must belong to the same schema as the view to be created.
- You must have `EXECUTE ANY TYPE` system privileges.
- You must have the `EXECUTE` object privilege on that object type.

## Partition Views

Partition views were introduced in Release 7.3 to provide partitioning capabilities for applications requiring them. Partition views are supported in Oracle8i so that you can upgrade applications from Release 7.3 without any modification. In most cases, subsequent to migration to Oracle8i you will want to migrate partition views into partitions (see *Oracle8i Administrator's Guide*).

With Oracle8i, you can use the `CREATE TABLE` statement to create partitioned tables easily. Partitioned tables offer the same advantages as partition views, while also addressing their shortcomings. Oracle recommends that you use partitioned tables rather than partition views in most operational environments.

**See Also:**

- *Oracle8i Concepts* for more information on the shortcomings of partition views.
- ["CREATE TABLE"](#) on page 4-381 for more information about partitioned tables.

**Keywords and Parameters**

---

OR REPLACE	<p>re-creates the view if it already exists. You can use this clause to change the definition of an existing view without dropping, re-creating, and regranted object privileges previously granted on it.</p> <p>INSTEAD OF triggers defined in the view are dropped when a view is re-created. See <a href="#">"CREATE TRIGGER"</a> on page 7-426 for more information about the INSTEAD OF clause.</p> <hr/> <p><b>Note:</b> If any materialized views are dependent on <i>view</i>, those materialized views will be marked UNUSABLE and will require a full refresh to restore them to a usable state. Invalid materialized views cannot be used by query rewrite and cannot be refreshed until they are recompiled. For information on refreshing invalid materialized views, see <a href="#">"ALTER MATERIALIZED VIEW / SNAPSHOT"</a> on page 7-47. For information on materialized views in general, see <i>Oracle8i Concepts</i>.</p> <hr/>
FORCE	<p>creates the view regardless of whether the view's base tables or the referenced object types exist or the owner of the schema containing the view has privileges on them. These conditions must be true before any SELECT, INSERT, UPDATE, or DELETE statements can be issued against the view.</p>
NO FORCE	<p>creates the view only if the base tables exist and the owner of the schema containing the view has privileges on them. This is the default.</p>
<i>schema</i>	<p>is the schema to contain the view. If you omit <i>schema</i>, Oracle creates the view in your own schema.</p>
<i>view</i>	<p>is the name of the view or the object view.</p> <p><b>Restriction:</b> If a view has INSTEAD OF triggers, any views created on it must have INSTEAD OF triggers, even if the views are inherently updatable.</p>
<i>alias</i>	<p>specifies names for the expressions selected by the view's query. The number of aliases must match the number of expressions selected by the view. Aliases must follow the rules for naming schema objects in the section, <a href="#">"Referring to Schema Objects and Parts"</a> on page 2-76. Aliases must be unique within the view.</p> <p>If you omit the aliases, Oracle derives them from the columns or column aliases in the view's query. For this reason, you <b>must</b> use aliases if the view's query contains expressions rather than only column names.</p>

---

	<p><b>Restriction:</b> You cannot specify an alias when creating an object view.</p>
OF <i>type_name</i>	<p>explicitly creates an object view of type <i>type_name</i>. The columns of an object view correspond to the top-level attributes of type <i>type_name</i>. Each row will contain an object instance and each instance will be associated with an object identifier (OID) as specified in the WITH OBJECT IDENTIFIER clause. If you omit <i>schema</i>, Oracle creates the object view in your own schema.</p> <p><b>See Also:</b> "CREATE TYPE" on page 7-437 for more information about creating objects.</p>
WITH OBJECT IDENTIFIER	<p>specifies the attributes of the object type that will be used as a key to identify each row in the object view. In most cases these attributes correspond to the primary-key columns of the base table. You must ensure that the attribute list is unique and identifies exactly one row in the view.</p> <p>If you try to dereference or pin a primary key REF that resolves to more than one instance in the object view, Oracle raises an error.</p> <hr/> <p><b>Note:</b> The 8.0 syntax WITH OBJECT OID is replaced with this syntax for clarity. The keywords WITH OBJECT OID are supported for backward compatibility, but Oracle Corporation recommends that you use the new syntax WITH OBJECT IDENTIFIER.</p> <hr/> <p>If the object view is defined on an object table or an object view, you can omit this clause or specify DEFAULT.</p>
DEFAULT	<p>specifies that the intrinsic object identifier of the underlying object table or object view will be used to uniquely identify each row.</p>
<i>attribute</i>	<p>is an attribute of the object type from which the object identifier for the object view is to be created.</p>
AS <i>subquery</i>	<p>identifies columns and rows of the table(s) that the view is based on. The subquery's select list can contain up to 1000 expressions.</p> <p>If you create views that refer to remote tables and views, the database links you specify must have been created using the CONNECT TO clause of the CREATE DATABASE LINK statement, and you must qualify them with schema name in the view query.</p>

**Restrictions:**

- The view query cannot select the `CURRVAL` or `NEXTVAL` pseudocolumns.
- If the view query selects the `ROWID`, `ROWNUM`, or `LEVEL` pseudocolumns, those columns must have aliases in the view query.
- If the view query uses an asterisk (\*) to select all columns of a table, and you later add new columns to the table, the view will not contain those columns until you re-create the view by issuing a `CREATE OR REPLACE VIEW` statement.
- For object views, the number of elements in the view subquery select list must be the same as the number of top-level attributes for the object type. The datatype of each of the selecting elements must be the same as the corresponding top-level attribute.
- You cannot specify the `SAMPLE` clause.

The preceding restrictions apply to materialized views as well.

- If you want the view to be inherently updatable, it must not contain any of the following constructs:
  - A set operator
  - A `DISTINCT` operator
  - An aggregate function
  - A `GROUP BY`, `ORDER BY`, `CONNECT BY`, or `START WITH` clause
  - A collection expression in a `SELECT` list
  - A subquery in a `SELECT` list
  - Joins (with some exceptions). See *Oracle8i Administrator's Guide* for details.
- If an inherently updatable view contains pseudocolumns or expressions, the `UPDATE` statement must not refer to any of these pseudocolumns or expressions.
- If you want a join view to be updatable, all of the following conditions must be true:
  - The DML statement must affect only one table underlying the join.
  - For an `UPDATE` statement, all columns updated must be extracted from a key-preserved table. If the view has the `CHECK OPTION`, join columns and columns taken from tables that are referenced more than once in the view must be shielded from `UPDATE`.
  - For a `DELETE` statement, the join can have one and only one key-preserved table. That table can appear more than once in the join, unless the view has the `CHECK OPTION`.
  - For an `INSERT` statement, all columns into which values are inserted must come from a key-preserved table, and the view must not have the `CHECK OPTION`.

**See Also:**

- *Oracle8i Administrator's Guide* for more information on updatable views.
- *Oracle8i Application Developer's Guide - Fundamentals* for more information about updating object views or relational views that support object types.



---

<i>with_clause</i>	restricts the subquery in one of the following ways:
WITH READ ONLY	specifies that no delete, inserts, or updates can be performed through the view.
WITH CHECK OPTION	specifies that inserts and updates performed through the view must result in rows that the view query can select. The CHECK OPTION cannot make this guarantee if: <ul style="list-style-type: none"> <li>■ There is a subquery in the query of this view or any view on which this view is based or</li> <li>■ INSERT, UPDATE, or DELETE operations are performed using INSTEAD OF triggers.</li> </ul>
CONSTRAINT <i>constraint</i>	assigns the name of the CHECK OPTION constraint. If you omit this identifier, Oracle automatically assigns the constraint a name of the form SYS_Cn, where n is an integer that makes the constraint name unique within the database.

---

## Examples

**Basic View Example** The following statement creates a view of the EMP table named DEPT20. The view shows the employees in Department 20 and their annual salary:

```
CREATE VIEW dept20
  AS SELECT ename, sal*12 annual_salary
     FROM emp
     WHERE deptno = 20;
```

The view declaration need not define a name for the column based on the expression SAL\*12, because the subquery uses a column alias (ANNUAL\_SALARY) for this expression.

**Updatable View Example** The following statement creates an updatable view named CLERKS of all clerks in the EMP table. Only the employees' IDs, names, and department numbers are visible in this view and only these columns can be updated in rows identified as clerks:

```
CREATE VIEW clerk (id_number, person, department, position)
  AS SELECT empno, ename, deptno, job
     FROM emp
     WHERE job = 'CLERK'
  WITH CHECK OPTION CONSTRAINT wco;
```

Because of the `CHECK OPTION`, you cannot subsequently insert a new row into `CLERK` if the new employee is not a clerk.

**Join View Example** A join view is one whose view query contains a join. If at least one column in the join has a unique index, then it may be possible to modify one base table in a join view. You can query `USER_UPDATABLE_COLUMNS` to see whether the columns in a join view are updatable. For example:

```
CREATE VIEW ed AS
  SELECT e.empno, e.ename, d.deptno, d.loc
     FROM emp e, dept d
     WHERE e.deptno = d.deptno
```

View created.

```
SELECT column_name, updatable
   FROM user_updatable_columns
   WHERE table_name = 'ED';
```

COLUMN_NAME	UPD
-----	---
ENAME	YES
DEPTNO	NO
EMPNO	YES
LOC	NO

```
INSERT INTO ed (ENAME, EMPNO) values ('BROWN', 1234);
```

In the above example, there is a unique index on the `DEPTNO` column of the `DEPT` table. You can insert, update or delete a row from the `EMP` base table, because all the columns in the view mapping to the `EMP` table are marked as updatable and because the primary key of `EMP` is included in the view.

---

---

**Note:** You cannot insert into the table using the view unless the view contains all `NOT NULL` columns of all tables in the join, unless you have specified `DEFAULT` values for the `NOT NULL` columns.

---

---

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals* for more information on updating join views.

**Read-Only View Example** The following statement creates a read-only view named `CLERKS` of all clerks in the `EMP` table. Only the employees' IDs, names, department numbers, and jobs are visible in this view:

```
CREATE VIEW clerk (id_number, person, department, position)
  AS SELECT empno, ename, deptno, job
     FROM emp
     WHERE job = 'CLERK'
  WITH READ ONLY;
```


**Object View Example** The following example creates object view `EMP_OBJECT_VIEW` of `EMPLOYEE_TYPE`:

```
CREATE TYPE employee_type AS OBJECT
( empno      NUMBER(4),
  ename      VARCHAR2(20),
  job        VARCHAR2(9),
  mgr        NUMBER(4),
  hiredate   DATE,
  sal        NUMBER(7,2),
  comm       NUMBER(7,2) );

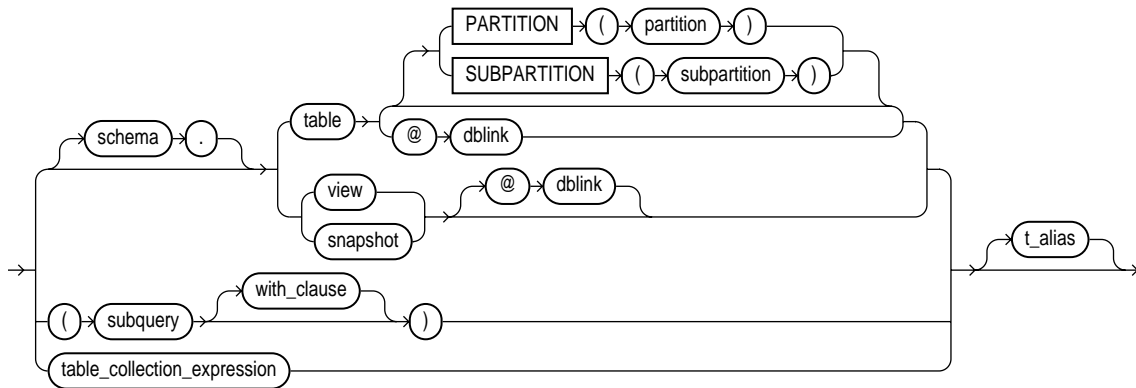
CREATE OR REPLACE VIEW emp_object_view OF employee_type
  WITH OBJECT IDENTIFIER (empno)
  AS SELECT empno, ename, job, mgr, hiredate, sal, comm
     FROM emp;
```

# DELETE

## Syntax




### DML\_table\_expression\_clause::=






**subquery:** See "SELECT and Subqueries" on page 7-569.

### with\_clause::=



### table\_collection\_expression::=



**where\_clause::=****returning\_clause::=****Purpose**

To remove rows from a table, a partitioned table, a view's base table, or a view's partitioned base table.

**Prerequisites**

For you to delete rows from a table, the table must be in your own schema or you must have `DELETE` privilege on the table.

For you to delete rows from the base table of a view, the owner of the schema containing the view must have `DELETE` privilege on the base table. Also, if the view is in a schema other than your own, you must be granted `DELETE` privilege on the view.

The `DELETE ANY TABLE` system privilege also allows you to delete rows from any table or table partition, or any view's base table.

If the `SQL92_SECURITY` initialization parameter is set to `TRUE`, then you must have `SELECT` privilege on the table to perform a `DELETE` that references table columns (such as the columns in a *where\_clause*).

**Keywords and Parameters**


---

<i>hint</i>	is a comment that passes instructions to the optimizer on choosing an execution plan for the statement. For the syntax and description of hints, see " <a href="#">Hints</a> " on page 2-63 and <i>Oracle8i Designing and Tuning for Performance</i> .
<i>DML_table_expression_clause</i>	
<i>schema</i>	is the schema containing the table or view. If you omit <i>schema</i> , Oracle assumes the table or view is in your own schema.

*table* | *view*  
/ *snapshot* |  
*subquery*

is the name of a table or view, or the column or columns resulting from a subquery, from which the rows are to be deleted. If you specify *view*, Oracle deletes rows from the view's base table.

If table (or the base table of view) contains one or more domain index columns, this statements executes the appropriate indextype delete routine.

**See Also:** *Oracle8i Data Cartridge Developer's Guide* for more information on these routines.

Issuing a DELETE statement against a table fires any DELETE triggers defined on the table.

All table or index space released by the deleted rows is retained by the table and index.

**Restrictions:**

- You cannot execute this statement if *table* (or the base table of *view*) contains any domain indexes marked LOADING or FAILED.
- You cannot specify the ORDER BY clause in the subquery of the *DML\_query\_expression\_clause*.
- You cannot delete from a view except through INSTEAD OF triggers if the view's defining query contains one of the following constructs:
  - A set operator
  - A DISTINCT operator
  - An aggregate function
  - A GROUP BY, ORDER BY, CONNECT BY, or START WITH clause
  - A collection expression in a SELECT list
  - A subquery in a SELECT list
  - Joins (with some exceptions). See *Oracle8i Administrator's Guide* for details.
- If you specify an index, index partition, or index subpartition that has been marked UNUSABLE, the DELETE statement will fail unless the SKIP\_UNUSABLE\_INDEXES parameter has been set to TRUE.

**See Also:** "ALTER SESSION" on page 7-83.

PARTITION  
(*partition\_name*) |  
SUBPARTITION  
(*subpartition\_*  
*name*)

specifies that *partition\_name* or *subpartition\_name* is the name of the partition or subpartition within *table* targeted for deletes.

You need not specify the partition name when deleting values from a partitioned table. However, in some cases, specifying the partition name is more efficient than a complicated *where\_clause*.

*dblink*

is the complete or partial name of a database link to a remote database where the table or view is located. For information on referring to database links, see "Referring to Objects in Remote Databases" on page 2-79. You can delete rows from a remote table or view only if you are using Oracle's distributed functionality.

If you omit *dblink*, Oracle assumes that the table or view is located on the local database.

---

<i>with_clause</i>	<p>restricts the subquery in one of the following ways:</p> <ul style="list-style-type: none"> <li>■ WITH READ ONLY specifies that the subquery cannot be updated.</li> <li>■ WITH CHECK OPTION specifies that Oracle prohibits any changes to that table that would produce rows that are not included in the subquery. See "WITH CHECK OPTION Example" on page 7-587.</li> </ul>
<i>table_collection_expression</i>	<p>informs Oracle that the collection value expression should be treated as a table. You can use a <i>table_collection_expression</i> to delete only those rows that also exist in another table.</p> <p><i>collection_expression</i> is a subquery that selects a nested table column from <i>table</i> or <i>view</i>.</p> <hr/> <p><b>Note:</b> In earlier releases of Oracle, <i>table_collection_expression</i> was expressed as "THE <i>subquery</i>". That usage is now deprecated.</p>
<i>where_clause</i>	<p>deletes only rows that satisfy the condition. The condition can reference the table and can contain a subquery. See the syntax description in "Conditions" on page 5-14. You can delete rows from a remote table or view only if you are using Oracle's distributed functionality.</p> <hr/> <p><b>Note:</b> If this clause contains a <i>subquery</i> that refers to remote objects, the DELETE operation can run in parallel as long as the reference does not loop back to an object on the local database. (See the description of the <i>parallel_clause</i> in "CREATE TABLE" on page 4-381.) However, if the <i>subquery</i> in the <i>DML_query_expression_clause</i> refers to any remote objects, the UPDATE operation will run serially without notification.</p> <hr/> <p>If you omit <i>dblink</i>, Oracle assumes that the table or view is located on the local database.</p> <p>If you omit the <i>where_clause</i>, Oracle deletes all rows of the table or view.</p>
<i>t_alias</i>	<p>provides a <b>correlation name</b> for the table, view, subquery, or collection value to be referenced elsewhere in the statement. Table aliases are generally used in DELETE statements with correlated queries.</p> <hr/> <p><b>Note:</b> This alias is required if the <i>DML_query_expression_clause</i> references any object type attributes or object type methods.</p>
<i>returning_clause</i>	<p>retrieves the rows affected by the DELETE statement.</p> <p>You can use a <i>returning_clause</i> to return values from deleted columns, and thereby eliminate the need to issue a SELECT statement following the DELETE statement.</p> <ul style="list-style-type: none"> <li>■ When deleting a single row, a DELETE statement with a <i>returning_clause</i> can retrieve column expressions using the deleted row, rowid, and REFs to the deleted row and store them in PL/SQL variables or bind variables.</li> <li>■ When deleting multiple rows, a DELETE statement with the <i>returning_clause</i> stores values from expressions, rowids, and REFs involving the deleted rows in bind arrays.</li> </ul> <p>You can also use DELETE with a <i>returning_clause</i> to delete from views with single base tables. For host binds, the datatype and size of the expression must be compatible with the bind variable.</p>

*expr* is any of the syntax descriptions in "Expressions" on page 5-1. You must specify a column expression in the *returning\_clause* for each variable in the *data\_item* list.

INTO indicates that the values of the changed rows are to be stored in the variable(s) specified in *data\_item* list.

*data\_item* is a PL/SQL variable or bind variable that stores the retrieved *expr* value.

**Restrictions:**

- You cannot use this clause with parallel DML or with remote objects.
  - You cannot retrieve LONG types with this clause.
  - You cannot specify this clause for a view on which an INSTEAD OF trigger has been defined.
- 

## Examples

**Basic Examples** The following statement deletes all rows from a table named TEMP\_ASSIGN.

```
DELETE FROM temp_assign;
```

The following statement deletes from the EMP table all sales staff who made less than \$100 commission last month:

```
DELETE FROM emp
  WHERE JOB = 'SALESMAN'
  AND COMM < 100;
```

The following statement has the same effect as the preceding example, but uses a subquery:

```
DELETE FROM (select * from emp)
  WHERE JOB = 'SALESMAN'
  AND COMM < 100;
```

**Remote Database Example** The following statement deletes all rows from the ACCOUNTS table owned by the user BLAKE on a database accessible by the database link DALLAS:

```
DELETE FROM blake.accounts@dallas;
```

**Nested Table Example** The following example deletes rows of nested table PROJS where the department number is either 123 or 456, or the department's budget is greater than 456.78:



---

```
DELETE THE(SELECT projs
            FROM dept d WHERE d.dno = 123) AS p
WHERE p.pno IN (123, 456) OR p.budgets > 456.78;
```

**Partition Example** The following example removes rows from partition NOV98 of the SALES table:

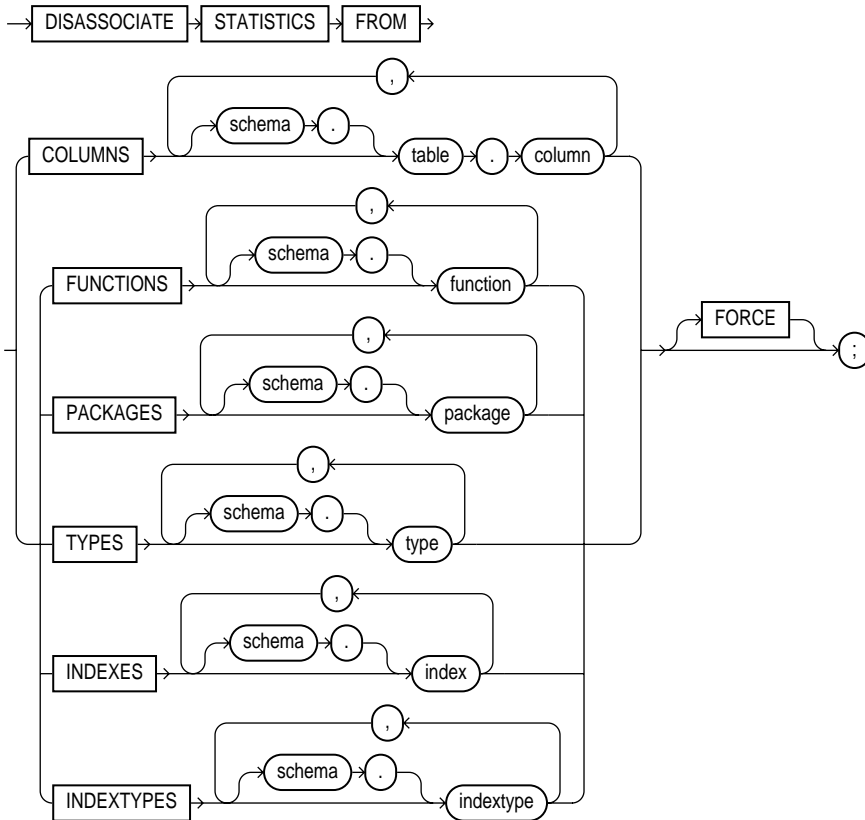
```
DELETE FROM sales PARTITION (nov98)
WHERE amount_of_sale != 0;
```

**Example** The following example returns column SAL from the deleted rows and stores the result in bind array :1:

```
DELETE FROM emp
WHERE job = 'SALESMAN' AND COMM < 100
RETURNING sal INTO :1;
```

## DISASSOCIATE STATISTICS

### Syntax



### Purpose

To disassociate a statistics type (or default statistics) from columns, standalone functions, packages, types, domain indexes, or indextypes.

**See Also:** "[ASSOCIATE STATISTICS](#)" on page 7-210 for more information on statistics type associations.

## Prerequisites

To issue this statement, you must have the appropriate privileges to alter the base object (table, function, package, type, domain index, or indextype).

## Keywords and Parameters

---

FROM COLUMNS	specifies a list of columns, standalone functions, packages, types, domain indexes, or indextypes from which you are disassociating statistics.
FUNCTIONS	
PACKAGES	If you do not specify <i>schema</i> , Oracle assumes the object is in your own schema.
TYPES	
INDEXES	If you have collected user-defined statistics on the object, the statement fails unless you specify <code>FORCE</code> .
INDEXTYPES	
FORCE	deletes the association regardless of whether any statistics exist for the object using the statistics type. If statistics do exist, the statistics are deleted before the association is deleted.

---

**Note:** When you drop an object with which a statistics type has been associated, Oracle automatically disassociates the statistics type with the `FORCE` option and drops all statistics that have been collected with the statistics type.

---

## Example

This statement disassociates statistics from the `PACK` package in the `HR` schema:

```
DISASSOCIATE STATISTICS FROM PACKAGES hr.pack;
```

## DROP CLUSTER

### Syntax



### Purpose

To remove a cluster from the database.

You cannot uncluster an individual table. Instead you must

- Create a new table with the same structure and contents as the old one, but with no `CLUSTER` clause,
- Drop the old table,
- Use the `RENAME` statement to give the new table the name of the old one, and
- Grant privileges on the new unclustered table, as grants on the old clustered table do not apply.

See ["CREATE TABLE"](#) on page 4-381, ["DROP TABLE"](#) on page 7-501, ["RENAME"](#) on page 7-554, ["GRANT system\\_privileges\\_and\\_roles"](#) on page 7-519.

### Prerequisites

The cluster must be in your own schema or you must have the `DROP ANY CLUSTER` system privilege.

### Keywords and Parameters

<i>schema</i>	is the schema containing the cluster. If you omit <i>schema</i> , Oracle assumes the cluster is in your own schema.
<i>cluster</i>	is the name of the cluster to be dropped. Dropping a cluster also drops the cluster index and returns all cluster space, including data blocks for the index, to the appropriate tablespace(s).

---

INCLUDING TABLES	drops all tables that belong to the cluster.
CASCADE CONSTRAINTS	drops all referential integrity constraints from tables outside the cluster that refer to primary and unique keys in tables of the cluster. If you omit this clause and such referential integrity constraints exist, Oracle returns an error and does not drop the cluster.

---

## Example

This statement drops a cluster named `GEOGRAPHY`, all its tables, and any referential integrity constraints that refer to primary or unique keys in those tables:

```
DROP CLUSTER geography  
  INCLUDING TABLES  
  CASCADE CONSTRAINTS;
```

## DROP CONTEXT

---

### Syntax

`DROP` → `CONTEXT` → `namespace` → `:`

### Purpose

To remove a context namespace from the database.

---

---

**Note:** Removing a context namespace does not invalidate any context under that namespace that has been set for a user session. However, the context will be invalid the next time the user attempts to set that context.

---

---

**See Also:** ["CREATE CONTEXT"](#) on page 7-261 and *Oracle8i Concepts* for more information on contexts.

### Prerequisites

You must have the `DROP ANY CONTEXT` system privilege.

### Keywords and Parameters

---

<i>namespace</i>	is the name of the context namespace to drop. You cannot drop the built-in namespace <code>USERENV</code> .
------------------	---

---

### Example

The following statement drops the context created in ["CREATE CONTEXT"](#) on page 7-261:

```
DROP CONTEXT hr_context;
```

---

## DROP DATABASE LINK

### Syntax



### Purpose

To remove a database link from the database.

### Prerequisites

To drop a private database link, the database link must be in your own schema. To drop a `PUBLIC` database link, you must have the `DROP PUBLIC DATABASE LINK` system privilege.

For information on creating database links, see ["CREATE DATABASE LINK"](#) on page 7-273.

### Keywords and Parameters

---

<code>PUBLIC</code>	must be specified to drop a <code>PUBLIC</code> database link.
<code>dblink</code>	specifies the database link to be dropped. <b>Restriction:</b> You cannot drop a database link in another user's schema and you cannot qualify <code>dblink</code> with the name of a schema. (Periods are permitted in names of database links. Therefore, Oracle interprets the entire name, such as <code>RALPH.LINKTOSALES</code> , as the name of a database link in your schema rather than as a database link named <code>LINKTOSALES</code> in the schema <code>RALPH</code> .)

---

### Example


The following statement drops a private database link named `BOSTON`:

```
DROP DATABASE LINK boston;
```

## DROP DIMENSION

---

### Syntax



### Purpose

To remove the named dimension.

For information on materialized views and their use of dimensions, see *Oracle8i Concepts*. See also "[CREATE DIMENSION](#)" on page 7-277.

### Prerequisites

The dimension must be in your own schema or you must have the `DROP ANY DIMENSION` system privilege to use this statement.

### Keywords and Parameters

---

<i>schema</i>	is the name of the schema in which the dimension is located. If you omit <i>schema</i> , Oracle assumes the dimension is in your own schema.
<i>dimension</i>	is the name of the dimension you want to drop. The dimension must already exist.  This statement does not invalidate materialized views that use relationships specified in dimension. However, requests that have been rewritten by query rewrite may be invalidated, and subsequent operations on such views may execute more slowly.

---

### Example

This example drops the `TIME` dimension:

```
DROP DIMENSION time;
```



---

## DROP DIRECTORY

### Syntax

```
DROP → DIRECTORY → directory_name → ;
```

### Purpose

Use `DROP DIRECTORY` to remove a directory object from the database.

For information on creating a directory, see "[CREATE DIRECTORY](#)" on page 7-282.

### Prerequisites

To drop a directory you must have the `DROP ANY DIRECTORY` system privilege.

---

---

**WARNING: Do not drop a directory when files in the associated file system are being accessed by PL/SQL or OCI programs.**

---

---

### Keywords and Parameters

---

<i>directory_name</i>	is the name of the directory database object to be dropped.
	Oracle removes the directory object, but does not delete the associated operating system directory on the server's file system.

---


### Example

The following statement drops the directory object `BFILE_DIR`:

```
DROP DIRECTORY bfile_dir;
```

## DROP FUNCTION

### Syntax



### Purpose

To remove a standalone stored function from the database. For information on creating a function, see ["CREATE FUNCTION"](#) on page 7-284.

---

---

**Note:** Do not use this statement to remove a function that is part of a package. Instead, either drop the entire package using the `DROP PACKAGE` statement or redefine the package without the function using the `CREATE PACKAGE` statement with the `OR REPLACE` clause.

---

---

### Prerequisites

The function must be in your own schema or you must have the `DROP ANY PROCEDURE` system privilege.

### Keywords and Parameters

---

*schema* is the schema containing the function. If you omit *schema*, Oracle assumes the function is in your own schema.

*function\_name* is the name of the function to be dropped.

Oracle invalidates any local objects that depend on, or call, the dropped function. If you subsequently reference one of these objects, Oracle tries to recompile the object and returns an error if you have not re-created the dropped function.

**See Also:** *Oracle8i Concepts* for more information on how Oracle maintains dependencies among schema objects, including remote objects.

---

If any statistics types are associated with the function, Oracle disassociates the statistics types with the `FORCE` option and drops any user-defined statistics collected with the statistics type.

**See Also:** "[ASSOCIATE STATISTICS](#)" on page 7-210 and "[DISASSOCIATE STATISTICS](#)" on page 7-470 for more information on statistics type associations.

---

## Example


The following statement drops the function `NEW_ACCT` in the schema `RIDDLEY` and invalidates all objects that depend upon `NEW_ACCT`:

```
DROP FUNCTION riddley.new_acct;
```

---

## DROP INDEX

### Syntax



### Purpose

To remove an index or domain index from the database.

#### See also:

- ["CREATE INDEX"](#) on page 7-291 and ["ALTER INDEX"](#) on page 7-30 for more information on indexes.
- The *domain\_index\_clause* of ["CREATE INDEX"](#) on page 7-291 for more information on domain indexes.

### Prerequisites

The index must be in your own schema or you must have the `DROP ANY INDEX` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the index. If you omit <i>schema</i> , Oracle assumes the index is in your own schema.
<i>index</i>	is the name of the index to be dropped. When the index is dropped, all data blocks allocated to the index are returned to the index's tablespace.
	If you drop a <b>domain index</b> :
	<ul style="list-style-type: none"> <li>■ Oracle invokes the appropriate indextype drop routine. For information on these routines, see <i>Oracle8i Data Cartridge Developer's Guide</i>.</li> <li>■ In addition, if any statistics are associated with the domain index, Oracle disassociates the statistics types with the <code>FORCE</code> clause and removes the user-defined statistics collected with the statistics type.</li> <li>■ <b>See Also:</b> <a href="#">"ASSOCIATE STATISTICS"</a> on page 7-210 and <a href="#">"DISASSOCIATE STATISTICS"</a> on page 7-470 for more information on statistics type associations.</li> </ul>

---

If you drop a global partitioned index, a range-partitioned, or a hash-partitioned index, all the index partitions are also dropped. If you drop a a composite-partitioned index, all the index partitions and subpartitions are also dropped.

FORCE

applies only to domain indexes. This clause drops the domain index even if the indextype routine invocation returns an error or the index is marked `LOADING`. Without `FORCE`, you cannot drop a domain index if its indextype routine invocation returns an error or the index is marked `LOADING`.

---


## Example

This statement drops an index named `MONOLITH`:

```
DROP INDEX monolith;
```

## DROP INDEXTYPE

### Syntax



### Purpose

To drop an indextype, as well as any association with a statistics type.

**See Also:** ["CREATE INDEXTYPE"](#) on page 7-309 for more information on indextypes.

### Prerequisites

The indextype must be in your own schema or you must have the DROP ANY INDEXTYPE system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the indextype. If you omit <i>schema</i> , Oracle assumes the indextype is in your own schema.
<i>indextype</i>	is the name of the indextype to be dropped.  If any statistics types have been associated with <i>indextype</i> , Oracle disassociates the statistics type from the indextype and drops any statistics that have been collected using the statistics type.  <b>See Also:</b> <a href="#">"ASSOCIATE STATISTICS"</a> on page 7-210 and <a href="#">"DISASSOCIATE STATISTICS"</a> on page 7-470 for more information on statistics associations.
FORCE	drops the indextype even if the indextype is currently being referenced by one or more domain indexes, and marks those domain indexes <code>INVALID</code> . Without <code>FORCE</code> , you cannot drop an indextype if any domain indexes reference the indextype.

---


### Example

The following statement drops the indextype `TEXTINDEXTYPE` and marks `INVALID` any domain indexes defined on this indextype:

```
DROP INDEXTYPE textindextype FORCE;
```

## DROP JAVA

### Syntax



### Purpose

To drop a Java source, class, or resource schema object.

**See Also:** *Oracle8i Java Stored Procedures Developer's Guide* for more information on resolving Java sources, classes, and resources.

### Prerequisites

The Java source, class, or resource must be in your own schema or you must have the `DROP ANY PROCEDURE` system privilege. You also must have the `EXECUTE` object privilege on Java classes to use this command.

### Keywords and Parameters

<code>JAVA SOURCE</code>	drops a Java source schema object and all Java class schema objects derived from it.
<code>JAVA CLASS</code>	drops a Java class schema object.
<code>JAVA RESOURCE</code>	drops a Java resource schema object.
<i>object_name</i>	specifies the name of an existing Java class, source, or resource schema object. Enclose the <i>object_name</i> in double quotation marks to preserve lower- or mixed-case names.

### Example

The following statement drops the Java class `MyClass`:

```
DROP JAVA CLASS "MyClass";
```

## DROP LIBRARY

---

### Syntax

`DROP` → `LIBRARY` → `library_name` → `:`

### Purpose

To remove an external procedure library from the database.

For information on creating a library, see "[CREATE LIBRARY](#)" on page 7-316.

### Prerequisites

You must have the `DROP LIBRARY` system privilege.

### Keywords and Parameters

---

<i>library_name</i>	is the name of the external procedure library being dropped.
---------------------	--

---

### Example


The following statement drops the `EXT_PROCS` library:

```
DROP LIBRARY ext_procs;
```



## DROP MATERIALIZED VIEW / SNAPSHOT

### Syntax



### Purpose

To remove an existing materialized view from the database.

The terms "snapshot" and "materialized view" are synonymous.

#### See Also:

- ["CREATE MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-318 for more information on materialized views, including a description of the various types of materialized views.
- *Oracle8i Replication* for information on materialized views in a replication environment.
- *Oracle8i Data Warehousing Guide* for information on materialized views in a data warehousing environment.

### Prerequisites

The materialized view must be in your own schema or you must have the `DROP ANY MATERIALIZED VIEW` (or `DROP ANY SNAPSHOT`) system privilege. You must also have the privileges to drop the internal table, views, and index that Oracle uses to maintain the materialized view's data.

For information on these privileges, see ["DROP TABLE"](#) on page 7-501, ["DROP VIEW"](#) on page 7-511, and ["DROP INDEX"](#) on page 7-480.

### Keywords and Parameters

---

*schema* is the schema containing the materialized view. If you omit *schema*, Oracle assumes the materialized view is in your own schema.

*materialized view* / *snapshot* is the name of the existing materialized view to be dropped.

- If you drop a simple materialized view that is the least recently refreshed materialized view of a master table, Oracle automatically purges from the detail table's materialized view log only the rows needed to refresh the dropped materialized view.
  - If you drop a detail table, Oracle does not automatically drop materialized views based on the table. However, Oracle returns an error when it tries to refresh a materialized view based on a detail table that has been dropped.
  - If you drop a materialized view, any compiled requests that were rewritten to use the materialized view will be invalidated and recompiled automatically. If the materialized view was prebuilt on a table, the table is not dropped, but it can no longer be maintained by the materialized view refresh mechanism.
- 

### Examples

The following statement drops the materialized view `PARTS` owned by the user `HQ`:


```
DROP SNAPSHOT hq.parts;
```

The following statement drops the `SALES_BY_MONTH` materialized view and the underlying table of the materialized view (unless the underlying table was registered in the `CREATE MATERIALIZED VIEW` statement with the `ON PREBUILT TABLE` clause):

```
DROP MATERIALIZED VIEW sales_by_month;
```

## DROP MATERIALIZED VIEW LOG / SNAPSHOT LOG

### Syntax



### Purpose

To remove a materialized view log from the database.

The terms "snapshot" and "materialized view" are synonymous.

#### See Also:

- ["CREATE MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-318 and ["ALTER MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-47 for more information on materialized views, including a description of the various types of materialized views and refreshing materialized views.
- ["CREATE MATERIALIZED VIEW LOG / SNAPSHOT LOG"](#) on page 7-333 for information on materialized view logs.
- *Oracle8i Replication* for information on materialized views in a replication environment.
- *Oracle8i Data Warehousing Guide* for information on materialized views in a data warehousing environment.

### Prerequisites

A materialized view log consists of a table and a trigger. To drop a materialized view log, you must have the privileges listed for ["DROP TABLE"](#) on page 7-501.

### Keywords and Parameters

---

*schema* is the schema containing the materialized view log and its master table. If you omit *schema*, Oracle assumes the materialized view log and master table are in your own schema.

---

*table* is the name of the detail table associated with the materialized view log to be dropped.

After you drop a materialized view log, some materialized views based on the materialized view log's detail table can no longer be fast refreshed. These materialized views include rowid materialized views, primary key materialized views, and subquery materialized views. For a description of the types of materialized views, see *Oracle8i Data Warehousing Guide*.

---

## Example


The following statement drops the materialized view log on the PARTS master table:

```
DROP MATERIALIZED VIEW LOG ON parts;
```

---

## DROP OPERATOR

### Syntax



### Purpose

To drop a user-defined operator.

**See Also:** ["User-Defined Operators"](#) on page 3-16, *Oracle8i Data Cartridge Developer's Guide* and ["CREATE OPERATOR"](#) on page 7-339 for more information on operators.

### Prerequisites

The operator must be in your schema or you must have the `DROP ANY OPERATOR` system privilege.

### Keywords and Parameters

<i>schema</i>	is the schema containing the operator. If you omit <i>schema</i> , Oracle assumes the operator is in your own schema.
<i>operator</i>	specifies the name of the operator to be dropped.
FORCE	drops the operator even if it is currently being referenced by one or more schema objects (indextypes, packages, functions, procedures, and so on), and marks those dependent objects <code>INVALID</code> . Without <code>FORCE</code> , you cannot drop an operator if any schema objects reference it.

### Example

The following statement drops the operator `MERGE`:

```
DROP OPERATOR ordsys.merge;
```

Because the `FORCE` clause is not specified, this operation will fail if any of the bindings of this operator are referenced by an indextype.

## DROP OUTLINE

---

### Syntax

```
DROP OUTLINE outline ;
```

### Purpose

To drop a stored outline.

**See Also:** ["CREATE OUTLINE"](#) on page 7-342 and *Oracle8i Designing and Tuning for Performance* for more information on outlines.

### Prerequisites

To drop an outline, you must have the `DROP ANY OUTLINE` system privilege.

### Keywords and Parameters

---

<i>outline</i>	is the name of the outline to be dropped. After the outline is dropped, if the SQL statement for which the stored outline was created is compiled, the optimizer generates a new execution plan without the influence of the outline.
----------------	--

---

### Example


The following statement drops the stored outline called `SALARIES`.

```
DROP OUTLINE salaries;
```

---

## DROP PACKAGE

### Syntax



### Purpose

To remove a stored package from the database. This statement drops the body and specification of a package.

---



---

**Note:** Do not use this statement to remove a single object from a package. Instead, re-create the package without the object using the `CREATE PACKAGE` and `CREATE PACKAGE BODY` statements with the `OR REPLACE` clause. See "[CREATE PACKAGE](#)" on page 7-344.

---



---

### Prerequisites

The package must be in your own schema or you must have the `DROP ANY PROCEDURE` system privilege.

### Keywords and Parameters

---

<code>BODY</code>	<p>drops only the body of the package. If you omit this clause, Oracle drops both the body and specification of the package.</p> <p>When you drop only the body of a package but not its specification, Oracle does not invalidate dependent objects. However, you cannot call one of the procedures or stored functions declared in the package specification until you re-create the package body.</p>
<code>schema</code>	<p>is the schema containing the package. If you omit <code>schema</code>, Oracle assumes the package is in your own schema.</p>

---

*package* is the name of the package to be dropped.

Oracle invalidates any local objects that depend on the package specification. If you subsequently reference one of these objects, Oracle tries to recompile the object and returns an error if you have not re-created the dropped package. For information on how Oracle maintains dependencies among schema objects, including remote objects, see *Oracle8i Concepts*.

If any statistics types are associated with the package, Oracle disassociates the statistics types with the FORCE clause and drops any user-defined statistics collected with the statistics types.

**See Also:** "[ASSOCIATE STATISTICS](#)" on page 7-210 and "[DISASSOCIATE STATISTICS](#)" on page 7-470.

---

### Example

The following statement drops the specification and body of the `BANKING` package, invalidating all objects that depend on the specification:


```
DROP PACKAGE banking;
```



---

## DROP PROCEDURE

### Syntax



### Purpose

To remove a standalone stored procedure from the database. Do not use this statement to remove a procedure that is part of a package. Instead, either drop the entire package using the `DROP PACKAGE` statement, or redefine the package without the procedure using the `CREATE PACKAGE` statement with the `OR REPLACE` clause.

For information on creating a procedure, see ["CREATE PROCEDURE"](#) on page 7-353.

### Prerequisites

The procedure must be in your own schema or you must have the `DROP ANY PROCEDURE` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the procedure. If you omit <i>schema</i> , Oracle assumes the procedure is in your own schema.
<i>procedure</i>	is the name of the procedure to be dropped.  When you drop a procedure, Oracle invalidates any local objects that depend upon the dropped procedure. If you subsequently reference one of these objects, Oracle tries to recompile the object and returns an error message if you have not re-created the dropped procedure. For information on how Oracle maintains dependencies among schema objects, including remote objects, see <i>Oracle8i Concepts</i> .

---

### Example

The following statement drops the procedure `TRANSFER` owned by the user `KERNER` and invalidates all objects that depend upon `TRANSFER`:

```
DROP PROCEDURE kerner.transfer
```

---

## DROP PROFILE

### Syntax



### Purpose

To remove a profile from the database.

For information on creating a profile, see ["CREATE PROFILE"](#) on page 7-359.

### Prerequisites

You must have the DROP PROFILE system privilege.

### Keywords and Parameters

---

<i>profile</i>	is the name of the profile to be dropped. <b>Restriction:</b> You cannot drop the DEFAULT profile.
CASCADE	deassigns the profile from any users to whom it is assigned. Oracle automatically assigns the DEFAULT profile to such users. You must specify this clause to drop a profile that is currently assigned to users.

---

### Example

The following statement drops the profile ENGINEER:


```
DROP PROFILE engineer CASCADE;
```

Oracle drops the profile ENGINEER and assigns the DEFAULT profile to any users currently assigned the ENGINEER profile.

---

## DROP ROLE

### Syntax



### Purpose

To remove a role from the database. When you drop a role, Oracle revokes it from all users and roles to whom it has been granted and removes it from the database.

For information on creating roles, see "[CREATE ROLE](#)" on page 4-365. For information on disabling roles for the current session, see "[SET ROLE](#)" on page 7-600.

### Prerequisites

You must have been granted the role with the `ADMIN OPTION` or you must have the `DROP ANY ROLE` system privilege.

### Keywords and Parameters

---

<i>role</i>	is the role to be dropped.
-------------	----------------------------

---

### Example

To drop the role `FLORIST`, issue the following statement:

```
DROP ROLE florist;
```

## DROP ROLLBACK SEGMENT

### Syntax

```
DROP → ROLLBACK → SEGMENT → rollback_segment → ;
```

### Purpose

To remove a rollback segment from the database. When you drop a rollback segment, all space allocated to the rollback segment returns to the tablespace.

For information on creating a rollback segment, see [CREATE ROLLBACK SEGMENT](#) on page 4-367. See also "[CREATE TABLESPACE](#)" on page 4-419.

### Prerequisites

You must have the `DROP ROLLBACK SEGMENT` system privilege.

### Keywords and Parameters

---

*rollback\_segment* is the name the rollback segment to be dropped.

**Restrictions:**

- You can drop a rollback segment only if it is offline. To determine whether a rollback segment is offline, query the data dictionary view `DBA_ROLLBACK_SEGS`. Offline rollback segments have the value `AVAILABLE` in the `STATUS` column. You can take a rollback segment offline with the `OFFLINE` clause of the [ALTER ROLLBACK SEGMENT](#) statement.
  - You cannot drop the `SYSTEM` rollback segment.
- 

### Example


The following statement drops the rollback segment `ACCOUNTING`:

```
DROP ROLLBACK SEGMENT accounting;
```

---

## DROP SEQUENCE

### Syntax



### Purpose

To remove a sequence from the database.

You can also use this statement to restart a sequence by dropping and then re-creating it. For example, if you have a sequence with a current value of 150 and you would like to restart the sequence with a value of 27, you can drop the sequence and then re-create it with the same name and a `START WITH` value of 27.

**See Also:** ["CREATE SEQUENCE"](#) on page 4-371 and ["ALTER SEQUENCE"](#) on page 7-81 for more information on creating and modifying sequences.

### Prerequisites

The sequence must be in your own schema or you must have the `DROP ANY SEQUENCE` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the sequence. If you omit <i>schema</i> , Oracle assumes the sequence is in your own schema.
<i>sequence_name</i>	is the name of the sequence to be dropped.

---

### Example

The following statement drops the sequence `ESEQ` owned by the user `ELLY`. To issue this statement, you must either be connected as user `ELLY` or have `DROP ANY SEQUENCE` system privilege:

```
DROP SEQUENCE elly.eseq;
```

---

## DROP SNAPSHOT

In Oracle8i, "snapshots" are synonymous with "materialized views." Please see ["DROP MATERIALIZED VIEW / SNAPSHOT"](#) on page 7-485.


## DROP SNAPSHOT LOG

---

In Oracle8i, "snapshots" are synonymous with "materialized views." Please see ["DROP MATERIALIZED VIEW LOG / SNAPSHOT LOG"](#) on page 7-487.

## DROP SYNONYM

### Syntax



### Purpose

To remove a synonym from the database, or to change the definition of a synonym by dropping and re-creating it.

**See Also:** ["CREATE SYNONYM"](#) on page 4-377 for more information on synonyms.

### Prerequisites

To drop a private synonym, either the synonym must be in your own schema or you must have the `DROP ANY SYNONYM` system privilege.

To drop a `PUBLIC` synonym, either the synonym must be in your own schema or you must have the `DROP ANY PUBLIC SYNONYM` system privilege.

### Keywords and Parameters

<code>PUBLIC</code>	must be specified to drop a public synonym. You cannot specify <i>schema</i> if you have specified <code>PUBLIC</code> .
<i>schema</i>	is the schema containing the synonym. If you omit <i>schema</i> , Oracle assumes the synonym is in your own schema.
<i>synonym</i>	is the name of the synonym to be dropped. If you drop a synonym for a materialized view, or its containing table or snapshot, or any of its dependent tables, the materialized view will be invalidated.

### Example


To drop a synonym named `MARKET`, issue the following statement:

```
DROP SYNONYM market;
```



## DROP TABLE

### Syntax



### Purpose

To remove a table or an object table and all its data from the database.

For information on creating tables, see ["CREATE TABLE"](#) on page 4-381. For information on modifying tables, see ["ALTER TABLE"](#) on page 7-123.

### Prerequisites

The table must be in your own schema or you must have the `DROP ANY TABLE` system privilege.

### Keywords and Parameters

<i>schema</i>	is the schema containing the table. If you omit <i>schema</i> , Oracle assumes the table is in your own schema.
<i>table</i>	is the name of the table, object table, or index-organized table to be dropped. Oracle automatically performs the following operations: <ul style="list-style-type: none"> <li>■ Removes all rows from the table (as if the rows were deleted).</li> <li>■ Drops all the table's indexes and domain indexes, regardless of who created them or whose schema contains them.</li> <li>■ If you drop a range-partitioned or hash-partitioned table, all the table partitions are also dropped. If you drop a composite-partitioned table, all the partitions and subpartitions are also dropped.</li> <li>■ For a domain index, this statement invokes the appropriate drop routines.</li> </ul> <p><b>See Also:</b> <i>Oracle8i Data Cartridge Developer's Guide</i> for more information on these routines.</p>

- If any statistic types are associated with the table, Oracle disassociates the statistics types with the `FORCE` clause and removes any user-defined statistics collected with the statistics type.

**See Also:** "[ASSOCIATE STATISTICS](#)" on page 7-210 and "[DISASSOCIATE STATISTICS](#)" on page 7-470 for more information on statistics type associations.

- If the table is not part of a cluster, Oracle returns all data blocks allocated to the table and its indexes to the tablespaces containing the table and its indexes.
- If the table is a base table for a view, a container or master table of a materialized view, or if it is referenced in a stored procedure, function, or package, Oracle invalidates these dependent objects but does not drop them. You cannot use these objects unless you re-create the table or drop and re-create the objects so that they no longer depend on the table.
- If you choose to re-create the table, it must contain all the columns selected by the queries originally used to define the materialized views/snapshots and all the columns referenced in the stored procedures, functions, or packages. Any users previously granted object privileges on the views, stored procedures, functions, or packages need not be regranted these privileges.
- If the table is a detail table for a materialized view, the materialized view can still be queried, but it cannot be refreshed unless the table is re-created so that it contains all the columns selected by the materialized view's query.
- If the table has a materialized view log/snapshot log, Oracle drops this log and any other direct-load `INSERT` refresh information associated with the table.

---

**Note:** To drop a cluster and all its the tables, use the `DROP CLUSTER` statement with the `INCLUDING TABLES` clause to avoid dropping each table individually. See "[DROP CLUSTER](#)" on page 7-472.

---

CASCADE  
CONSTRAINTS

drops all referential integrity constraints that refer to primary and unique keys in the dropped table. If you omit this clause, and such referential integrity constraints exist, Oracle returns an error and does not drop the table.

---


## Example

The following statement drops the `TEST_DATA` table:

```
DROP TABLE test_data;
```

## DROP TABLESPACE

### Syntax



### Purpose

To remove a tablespace from the database.

For information on creating and modifying tablespaces, see ["CREATE TABLESPACE"](#) on page 4-419 and ["ALTER TABLESPACE"](#) on page 7-176.

### Prerequisites

You must have the `DROP TABLESPACE` system privilege. You cannot drop a tablespace if it contains any rollback segments holding active transactions.

### Keywords and parameters

---

<i>tablespace</i>	<p>is the name of the tablespace to be dropped.</p> <p>You can drop a tablespace regardless of whether it is online or offline. Oracle recommends that you take the tablespace offline before dropping it to ensure that no SQL statements in currently running transactions access any of the objects in the tablespace.</p> <p>You may want to alert any users who have been assigned the tablespace as either a default or temporary tablespace. After the tablespace has been dropped, these users cannot allocate space for objects or sort areas in the tablespace. You can reassign users new default and temporary tablespaces with the <code>ALTER USER</code> statement.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ You cannot drop the <code>SYSTEM</code> tablespace.</li> <li>■ You cannot drop a tablespace that contains a domain index or any objects created by a domain index.</li> </ul> <p><b>See Also:</b> <i>Oracle8i Data Cartridge Developer's Guide</i> and <i>Oracle8i Concepts</i> for more information on domain indexes.</p>
-------------------	---

---

### INCLUDING CONTENTS

drops all the contents of the tablespace. You must specify this clause to drop a tablespace that contains any database objects. If you omit this clause, and the tablespace is not empty, Oracle returns an error and does not drop the tablespace.

For **partitioned tables**, DROP TABLESPACE will fail even if you specify INCLUDING CONTENTS, if the tablespace contains some, but not all,

- partitions of a range- or hash-partitioned table, or
- subpartitions of a composite-partitioned table.

**Note:** If all the partitions of a partitioned table reside in *tablespace*, DROP TABLESPACE ... INCLUDING CONTENTS will drop *tablespace*, as well as any associated index segments, LOB data segments, and LOB index segments in the other tablespace(s).

For a **partitioned index-organized table**, if all the primary key index segments are in this tablespace, this clause will also drop any overflow segments that exist in other tablespaces. If some of the primary key index segments are *not* in this tablespace, the statement will fail. In that case, before you can drop the tablespace, you must use ALTER TABLE ... MOVE PARTITION to move those primary key index segments into this tablespace, drop the partitions whose overflow data segments are not in this tablespace, and drop the partitioned index-organized table.

If the tablespace contains a container table or detail table of a materialized view, Oracle invalidates the materialized view.

If the tablespace contains a materialized view/snapshot log, Oracle drops this log and any other direct-load INSERT refresh information associated with the table.

### CASCADE CONSTRAINTS

drops all referential integrity constraints from tables outside *tablespace* that refer to primary and unique keys of tables inside *tablespace*. If you omit this clause and such referential integrity constraints exist, Oracle returns an error and does not drop the tablespace.

---

## Example


The following statement drops the MFRG tablespace and all its contents:

```
DROP TABLESPACE mfrg
  INCLUDING CONTENTS
  CASCADE CONSTRAINTS;
```

---

## DROP TRIGGER

### Syntax



### Purpose

To remove a database trigger from the database.

**See Also:** ["CREATE TRIGGER"](#) on page 7-426 for information on creating triggers.

### Prerequisites

The trigger must be in your own schema or you must have the `DROP ANY TRIGGER` system privilege.

In addition, to drop a trigger on `DATABASE` in another user's schema, you must have the `ADMINISTER DATABASE TRIGGER` system privilege.

**See Also:** ["CREATE TRIGGER"](#) on page 7-426 for information on these privileges.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the trigger. If you omit <i>schema</i> , Oracle assumes the trigger is in your own schema.
<i>trigger</i>	is the name of the trigger to be dropped. Oracle removes it from the database and does not fire it again.

---

### Example


The following statement drops the `REORDER` trigger in the schema `RUTH`:

```
DROP TRIGGER ruth.reorder;
```

---

## DROP TYPE

### Syntax



### Purpose

To drop the specification and body of an object, a varray, or nested table type. To drop just the body of an object type, see ["DROP TYPE BODY"](#) on page 7-508.

**See Also:** ["CREATE TYPE"](#) on page 7-437 for information on creating types.

### Prerequisites

The object, varray, or nested table type must be in your own schema or you must have the `DROP ANY TYPE` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the type. If you omit <i>schema</i> , Oracle assumes the type is in your own schema.
<i>type_name</i>	<p>is the name of the object, varray, or nested table type to be dropped. You can drop only types with no type or table dependencies.</p> <p>If <i>type_name</i> is a statistics type, this statement will fail unless you also specify <code>FORCE</code>. If you specify <code>FORCE</code>, Oracle first disassociates all objects that are associated with <i>type_name</i>, and then drops <i>type_name</i>.</p> <p><b>See Also:</b> <a href="#">"ASSOCIATE STATISTICS"</a> on page 7-210 and <a href="#">"DISASSOCIATE STATISTICS"</a> on page 7-470 for more information on statistics types.</p> <p>If <i>type_name</i> is an object type that has been associated with a statistics type, Oracle first attempts to disassociate <i>type_name</i> from the statistics type and then drop <i>type_name</i>. However, if statistics have been collected using the statistics type, Oracle will be unable to disassociate <i>type_name</i> from the statistics type, and this statement will fail.</p>

---

If *type\_name* is an implementation type for an indextype, the indextype will be marked `INVALID`.

Unless you specify `FORCE`, you can drop only object, nested table, or varray types that are standalone schema objects with no dependencies. This is the default behavior.

**See Also:** "[CREATE INDEXTYPE](#)" on page 7-309.

`FORCE`

forces the type to be dropped even if it has dependent database objects. Oracle marks `UNUSED` all columns dependent on the type to be dropped, and those columns become inaccessible.

---

**WARNING:** Oracle does not recommend that you specify `FORCE` to drop types with dependencies. This operation is not recoverable and could cause the data in the dependent tables or columns to become inaccessible. For information about type dependencies, see *Oracle8i Application Developer's Guide - Fundamentals*.

---

## Example


The following statement removes object type `PERSON_T`:

```
DROP TYPE person_t;
```

---

## DROP TYPE BODY

### Syntax



### Purpose

To drop the body of an object, varray, or nested table type. When you drop a type body, the object type specification still exists, and you can re-create the type body. Prior to re-creating the body, you can still use the object type, although you cannot call the member functions.

#### See Also:

- ["DROP TYPE"](#) on page 7-506 for information on dropping the specification of an object.
- ["CREATE TYPE BODY"](#) on page 7-447 for more information on type bodies.

### Prerequisites

The object type body must be in your own schema, and you must have

- the `CREATE TYPE` or `CREATE ANY TYPE` system privilege, or
- the `DROP ANY TYPE` system privilege

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the object type. If you omit <i>schema</i> , Oracle assumes the object type is in your own schema.
<i>type_name</i>	is the name of the object type body to be dropped. <b>Restriction:</b> You can drop a type body only if it has no type or table dependencies.

---

### Example

The following statement removes object type body `RATIONAL`:

```
DROP TYPE BODY rational;
```



---

## DROP USER

### Syntax



### Purpose

To remove a database user and optionally remove the user's objects. For information on creating a user, see "[CREATE USER](#)" on page 7-451. For information on modifying the definition of a user, see "[ALTER USER](#)" on page 7-193.

### Prerequisites

You must have the DROP USER system privilege.

### Keywords and Parameters

---

<i>user</i>	is the user to be dropped. Oracle does not drop users whose schemas contain objects unless you specify <i>CASCADE</i> , or unless you first explicitly drop the user's objects.
-------------	---

---

### CASCADE

drops all objects in the user's schema before dropping the user. You must specify this clause to drop a user whose schema contains any objects.

- If the user's schema contains tables, Oracle drops the tables and automatically drops any referential integrity constraints on tables in other schemas that refer to primary and unique keys on these tables.
- If this clause results in tables being dropped, Oracle also drops all domain indexes created on columns of those tables, and invokes appropriate drop routines.

**See Also:** *Oracle8i Data Cartridge Developer's Guide* for more information on these routines.

- Oracle invalidates, but does not drop, the following objects in **other** schemas: views or synonyms for objects in the dropped user's schema; and stored procedures, functions, or packages that query objects in the dropped user's schema.
- Oracle does not drop materialized views on tables or views in the dropped user's schema, but if you specify *CASCADE*, the materialized views can no longer be refreshed.
- Oracle drops all triggers in the user's schema.
- Oracle does **not** drop roles created by the user.

**WARNING:** Oracle also drops with *FORCE* all types owned by the user. See the **FORCE** keyword of "**DROP TYPE**" on page 7-507.

---

## Examples

If user *BRADLEY*'s schema contains no objects, you can drop *BRADLEY* by issuing the statement:

```
DROP USER bradley;
```


If *BRADLEY*'s schema contains objects, you must use the *CASCADE* clause to drop *BRADLEY* and the objects:

```
DROP USER bradley CASCADE;
```

---

## DROP VIEW

### Syntax



### Purpose

To remove a view or an object view from the database. You can change the definition of a view by dropping and re-creating it.

**See Also:** ["CREATE VIEW"](#) on page 7-456.

### Prerequisites

The view must be in your own schema or you must have the `DROP ANY VIEW` system privilege.

### Keywords and Parameters

---

<i>schema</i>	is the schema containing the view. If you omit <i>schema</i> , Oracle assumes the view is in your own schema.
<i>view</i>	is the name of the view to be dropped.  Views, materialized views, and synonyms that refer to the view are not dropped, but become invalid. You can drop them or redefine views and synonyms, or you can define other views in such a way that the invalid views and synonyms become valid again. See <a href="#">"CREATE TABLE"</a> on page 4-381 and <a href="#">"CREATE SYNONYM"</a> on page 4-377. To revalidate invalid materialized views, see <a href="#">"ALTER MATERIALIZED VIEW / SNAPSHOT"</a> on page 7-47.

---

### Example


The following statement drops the `VIEW_DATA` view:

```
DROP VIEW view_data;
```

---

## EXPLAIN PLAN

### Syntax



### Purpose

To determine the execution plan Oracle follows to execute a specified SQL statement. This statement inserts a row describing each step of the execution plan into a specified table. If you are using cost-based optimization, this statement also determines the cost of executing the statement. If any domain indexes are defined on the table, user-defined CPU and I/O costs will also be inserted. See *Oracle8i Designing and Tuning for Performance* for information on the output of EXPLAIN PLAN.

The definition of a sample output table `PLAN_TABLE` is available in a SQL script on your distribution media. Your output table must have the same column names and datatypes as this table. The common name of this script is `UTLXPLAN.SQL`. The exact name and location depend on your operating system.

---

**Note:** Do not use the EXPLAIN PLAN statement to determine the execution plans of SQL statements that access data dictionary views or dynamic performance tables.

---

You can also issue the EXPLAIN PLAN statement as part of the SQL trace facility. For information on how to use the SQL trace facility, as well as a detailed discussion of how to generate and interpret execution plans, see *Oracle8i Designing and Tuning for Performance*.

## Prerequisites

To issue an `EXPLAIN PLAN` statement, you must have the privileges necessary to insert rows into an existing output table that you specify to hold the execution plan. For information on these privileges, see ["INSERT"](#) on page 7-539.

You must also have the privileges necessary to execute the SQL statement for which you are determining the execution plan. If the SQL statement accesses a view, you must have privileges to access any tables and views on which the view is based. If the view is based on another view that is based on a table, you must have privileges to access both the other view and its underlying table.

To examine the execution plan produced by an `EXPLAIN PLAN` statement, you must have the privileges necessary to query the output table.

The `EXPLAIN PLAN` statement is a data manipulation language (DML) statement, rather than a data definition language (DDL) statement. Therefore, Oracle does not implicitly commit the changes made by an `EXPLAIN PLAN` statement. If you want to keep the rows generated by an `EXPLAIN PLAN` statement in the output table, you must commit the transaction containing the statement.

**See Also:** ["SELECT and Subqueries"](#) on page 7-569 for more information on these privileges.

## Keywords and Parameters

---

<code>SET STATEMENT_ID</code>	specifies the value of the <code>STATEMENT_ID</code> column for the rows of the execution plan in the output table. You can then use this value to identify these rows among others in the output table. Be sure to specify a <code>STATEMENT_ID</code> value if your output table contains rows from many execution plans. If you omit this clause, the <code>STATEMENT_ID</code> value defaults to null.
<code>INTO</code>	<p>specifies name of the output table, and optionally its schema and database. This table must exist before you use the <code>EXPLAIN PLAN</code> statement.</p> <p>If you omit <i>schema</i>, Oracle assumes the table is in your own schema.</p> <p>The <i>dblink</i> can be a complete or partial name of a database link to a remote Oracle database where the output table is located. For information on referring to database links, see the section, <a href="#">"Referring to Objects in Remote Databases"</a> on page 2-79. You can specify a remote output table only if you are using Oracle's distributed functionality. If you omit <i>dblink</i>, Oracle assumes the table is on your local database.</p> <p>If you omit <code>INTO</code> altogether, Oracle assumes an output table named <code>PLAN_TABLE</code> in your own schema on your local database.</p>

---

**FOR *statement*** specifies a SELECT, INSERT, UPDATE, DELETE, CREATE TABLE, or CREATE INDEX statement for which the execution plan is generated.

**Note:** If *statement* includes the *parallel\_clause*, the resulting execution plan will indicate parallel execution. However, EXPLAIN PLAN actually inserts the statement into the plan table, so that the parallel DML statement you submit is no longer the first DML statement in the transaction. This violates the Oracle restriction of one parallel DML statement per transaction, and the statement will be executed serially. To maintain parallel execution of the statements, you must commit or roll back the EXPLAIN PLAN statement, and then submit the parallel DML statement.

---

## Examples

The following statement determines the execution plan and cost for an UPDATE statement and inserts rows describing the execution plan into the specified OUTPUT table with the STATEMENT\_ID value of 'Raise in Chicago':

```
EXPLAIN PLAN
  SET STATEMENT_ID = 'Raise in Chicago'
  INTO output
  FOR UPDATE emp
    SET sal = sal * 1.10
    WHERE deptno = (SELECT deptno
                    FROM dept
                    WHERE loc = 'CHICAGO');
```

The following SELECT statement queries the OUTPUT table and returns the execution plan and the cost:

```
SELECT LPAD(' ', 2*(LEVEL-1)) || operation operation, options,
       object_name, position
  FROM output
  START WITH id = 0 AND statement_id = 'Raise in Chicago'
  CONNECT BY PRIOR id = parent_id AND
             statement_id = 'Raise in Chicago';
```

The query returns this execution plan:

OPERATION	OPTIONS	OBJECT_NAME	POSITION
UPDATE STATEMENT			1
FILTER			0
TABLE ACCESS	FULL	EMP	1
TABLE ACCESS	FULL	DEPT	2

The value in the `POSITION` column of the first row shows that the statement has a cost of 1.

**Partitioned Example** Assume that `STOCKS` is a table with eight partitions on a `STOCK_NUM` column, and that a local prefixed index `STOCK_IX` on column `STOCK_NUM` exists. The partition `HIGHVALUES` are 1000, 2000, 3000, 4000, 5000, 6000, 7000, and 8000.

Consider the query:

```
SELECT * FROM stocks WHERE stock_num BETWEEN 3800 AND :h;
```


(where `:h` represents a bind variable). `EXPLAIN PLAN` executes this query with `PLAN_TABLE` as the output table. The basic execution plan, including partitioning information, is obtained with the query:

```
SELECT id, operation, options, object_name,  
       partition_start, partition_stop, partition_id FROM plan_table;
```


## filespec

### Syntax

filespec\_datafiles & filespec\_tempfiles::=



filespec\_redo\_log\_file\_groups::=



### Purpose

To specify a file as a datafile or tempfile

To specify a group of one or more files as a redo log file group.

### Prerequisites

A *filespec* can appear in the following statements: ["CREATE DATABASE"](#) on page 7-267, ["ALTER DATABASE"](#) on page 7-6, ["CREATE TABLESPACE"](#) on page 4-419, and ["ALTER TABLESPACE"](#) on page 7-176, ["CREATE CONTROLFILE"](#) on page 7-263, ["CREATE LIBRARY"](#) on page 7-316, and ["CREATE TEMPORARY TABLESPACE"](#) on page 4-424.

You must have the privileges necessary to issue one of these statements.



## Keywords and Parameters

---

<i>'filename'</i>	<p>is the name of either a datafile, tempfile, or a redo log file member. A <i>'filename'</i> can contain only single-byte characters from 7-bit ASCII or EBCDIC character sets. Multibyte characters are not valid.</p> <p>A redo log file group can have one or more members (copies). Each <i>'filename'</i> must be fully specified according to the conventions for your operating system.</p>
SIZE <i>integer</i>	<p>specifies the size of the file. Use <b>K</b> or <b>M</b> to specify the size in kilobytes or megabytes.</p> <ul style="list-style-type: none"> <li>■ You can omit this parameter <b>only</b> if the file already exists.</li> <li>■ The size of a tablespace must be one block greater than the sum of the sizes of the objects contained in it.</li> </ul>
REUSE	<p>allows Oracle to reuse an existing file.</p> <ul style="list-style-type: none"> <li>■ If the file already exists, Oracle verifies that its size matches the value of the <b>SIZE</b> parameter (if you specify <b>SIZE</b>).</li> <li>■ If the file does not exist, Oracle ignores this clause and creates the file.</li> <li>■ You can omit this clause only if the file does not already exist. If you omit this clause, Oracle creates the file.</li> </ul>

---

**Note:** Whenever Oracle uses an existing file, the file's previous contents are lost.

---

## Examples

The following statement creates a database named **PAYABLE** that has two redo log file groups, each with two members, and one datafile:

```
CREATE DATABASE payable
  LOGFILE GROUP 1 ('diska:log1.log', 'diskb:log1.log') SIZE 50K,
  GROUP 2 ('diska:log2.log', 'diskb:log2.log') SIZE 50K
  DATAFILE 'diskc:dbone.dat' SIZE 30M;
```

The first *filespec* in the **LOGFILE** clause specifies a redo log file group with the **GROUP** value 1. This group has members named *'diska:log1.log'* and *'diskb:log1.log'*, each 50 kilobytes in size.

The second *filespec* in the **LOGFILE** clause specifies a redo log file group with the **GROUP** value 2. This group has members named *'diska:log2.log'* and *'diskb:log2.log'*, also 50 kilobytes in size.

The *filespec* in the **DATAFILE** clause specifies a datafile named *'diskc:dbone.dat'*, 30 megabytes in size.

All of these *filespecs* specify a value for the `SIZE` parameter and omit the `REUSE` clause, so none of these files can already exist. Oracle must create them.

The following statement adds another redo log file group with two members to the `PAYABLE` database:

```
ALTER DATABASE payable
  ADD LOGFILE GROUP 3 ('diska:log3.log', 'diskb:log3.log')
  SIZE 50K REUSE;
```

The *filespec* in the `ADD LOGFILE` clause specifies a new redo log file group with the `GROUP` value 3. This new group has members named `'diska:log3.log'` and `'diskb:log3.log'`, each 50 kilobytes in size. Because the *filespec* specifies the `REUSE` clause, each member can already exist. If a member exists, it must have a size of 50 kilobytes. If it does not exist, Oracle creates it with that size.

The following statement creates a tablespace named `STOCKS` that has three datafiles:

```
CREATE TABLESPACE stocks
  DATAFILE 'diskc:stock1.dat',
           'diskc:stock2.dat',
           'diskc:stock3.dat';
```

The *filespecs* for the datafiles specifies files named `'diskc:stock1.dat'`, `'diskc:stock2.dat'`, and `'diskc:stock3.dat'`. Since each *filespec* omits the `SIZE` parameter, each file must already exist.


The following statement alters the `STOCKS` tablespace and adds a new datafile:

```
ALTER TABLESPACE stocks
  ADD DATAFILE 'diskc:stock4.dat' REUSE;
```

The *filespec* specifies a datafile named `'diskc:stock4.dat'`. Since the *filespec* omits the `SIZE` parameter, the file must already exist and the `REUSE` clause is not significant.

## GRANT *system\_privileges\_and\_roles*

### Syntax



### Purpose

To grant system privileges and roles to users and roles. Both privileges and roles are either local, global, or external (see ["CREATE USER"](#) on page 7-451 and ["CREATE ROLE"](#) on page 4-365 for definitions).

You can authorize database users to use roles through means other than the database and the GRANT statement. For example, some operating systems have facilities that grant operating system privileges to operating system users. You can use such facilities to grant roles to Oracle users with the initialization parameter `OS_ROLES`. If you choose to grant roles to users through operating system facilities, you cannot also grant roles to users with the GRANT statement, although you can use the GRANT statement to grant system privileges to users and system privileges and roles to other roles. For information about other authorization methods, see *Oracle8i Administrator's Guide*.

For information on granting object privileges, see ["GRANT object\\_privileges"](#) on page 7-532.

### Prerequisites

To grant a system privilege, you must either have been granted the system privilege with the `ADMIN OPTION` or have been granted the `GRANT ANY PRIVILEGE` system privilege.

To grant a role, you must either have been granted the role with the `ADMIN OPTION` or have been granted the `GRANT ANY ROLE` system privilege, or you must have created the role.

## Keywords and Parameters

<i>system_priv</i>	<p>is a system privilege to be granted. <a href="#">Table 7-5</a> lists the system privileges (organized by the database object operated upon).</p> <ul style="list-style-type: none"> <li>■ If you grant <b>a privilege to a user</b>, Oracle adds the privilege to the user's privilege domain. The user can immediately exercise the privilege.</li> <li>■ If you grant <b>a privilege to a role</b>, Oracle adds the privilege to the role's privilege domain. Users who have been granted and have enabled the role can immediately exercise the privilege. Other users who have been granted the role can enable the role and exercise the privilege.</li> <li>■ If you grant <b>a privilege to PUBLIC</b>, Oracle adds the privilege to the privilege domains of each user. All users can immediately perform operations authorized by the privilege.</li> </ul> <p>Oracle provides a shortcut for specifying all system privileges at once:</p> <p><b>ALL PRIVILEGES</b> is equivalent to specifying all the system privileges listed in <a href="#">Table 7-5, "System Privileges"</a> on page 7-521.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ A privilege or role cannot appear more than once in the list of privileges and roles to be granted.</li> <li>■ You cannot grant a role to itself.</li> <li>■ You cannot grant a role <code>IDENTIFIED GLOBALLY</code> to anything.</li> <li>■ You cannot grant a role <code>IDENTIFIED EXTERNALLY</code> to a global user or global role.</li> <li>■ You cannot grant roles circularly. For example, if you grant the role <code>BANKER</code> to the role <code>TELLER</code>, you cannot subsequently grant <code>TELLER</code> to <code>BANKER</code>.</li> </ul>
<i>role</i>	<p>is a role to be granted. You can grant an Oracle predefined role or a user-defined role. <a href="#">Table 7-6</a> lists the predefined roles. For information on creating a user-defined role, see "<a href="#">CREATE ROLE</a>" on page 4-365.</p> <ul style="list-style-type: none"> <li>■ If you grant <b>a role to a user</b>, Oracle makes the role available to the user. The user can immediately enable the role and exercise the privileges in the role's privilege domain.</li> <li>■ If you grant <b>a role to another role</b>, Oracle adds the granted role's privilege domain to the grantee role's privilege domain. Users who have been granted the grantee role can enable it and exercise the privileges in the granted role's privilege domain.</li> <li>■ If you grant <b>a role to PUBLIC</b>, Oracle makes the role available to all users. All users can immediately enable the role and exercise the privileges in the roles privilege domain.</li> </ul>
TO	<p>identifies users or roles to which system privileges and roles are granted.</p> <p><b>Restriction:</b> A user, role, or <code>PUBLIC</code> cannot appear more than once in the <code>TO</code> clause.</p>
PUBLIC	<p>grants system privileges or roles to all users.</p>

---

WITH ADMIN  
OPTION

enables the grantee to

- Grant the role to another user or role, unless the role is a GLOBAL role
- Revoke the role from another user or role
- Alter the role to change the authorization needed to access it
- Drop the role

If you grant a system privilege or role to a user without specifying WITH ADMIN OPTION, and then subsequently grant the privilege or role to the user WITH ADMIN OPTION, the user has the ADMIN OPTION on the privilege or role.

To revoke the admin option on a system privilege or role from a user, you must revoke the privilege or role from the user altogether and then grant the privilege or role to the user without the admin option.

---

**Table 7-5 System Privileges**

System Privilege	Allows grantee to . . .
<b>CLUSTERS</b>	
CREATE CLUSTER	Create clusters in grantee's schema
CREATE ANY CLUSTER	Create a cluster in any schema. Behaves similarly to CREATE ANY TABLE.
ALTER ANY CLUSTER	Alter clusters in any schema
DROP ANY CLUSTER	Drop clusters in any schema
<b>CONTEXTS</b>	
CREATE ANY CONTEXT	Create any context namespace
DROP ANY CONTEXT	Drop any context namespace
<b>DATABASE</b>	
ALTER DATABASE	Alter the database
ALTER SYSTEM	Issue ALTER SYSTEM statements
AUDIT SYSTEM	Issue AUDIT <i>sql_statements</i> statements
<b>DATABASE LINKS</b>	

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter O7\_DICTIONARY\_ACCESSIBILITY to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7-5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
CREATE DATABASE LINK	Create private database links in grantee's schema
CREATE PUBLIC DATABASE LINK	Create public database links
DROP PUBLIC DATABASE LINK	Drop public database links
<b>DIMENSIONS</b>	
CREATE DIMENSION	Create dimensions in the grantee's schema
CREATE ANY DIMENSION	Create dimensions in any schema
ALTER ANY DIMENSION	Alter dimensions in any schema
DROP ANY DIMENSION	Drop dimensions in any schema
<b>DIRECTORIES</b>	
CREATE ANY DIRECTORY	Create directory database objects
DROP ANY DIRECTORY	Drop directory database objects
<b>INDEXTYPES</b>	
CREATE INDEXTYPE	Create an indextype in the grantee's schema
CREATE ANY INDEXTYPE	Create an indextype in any schema
ALTER ANY INDEXTYPE	Modify indextypes in any schema
DROP ANY INDEXTYPE	Drop an indextype in any schema
EXECUTE ANY INDEXTYPE	Reference an indextype in any schema
<b>INDEXES</b>	
CREATE ANY INDEX	Create in any schema a domain index or an index on any table in any schema
ALTER ANY INDEX	Alter indexes in any schema
DROP ANY INDEX	Drop indexes in any schema

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter `O7_DICTIONARY_ACCESSIBILITY` to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7-5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
QUERY REWRITE	Enable rewrite using a materialized view, or create a function-based index, when that materialized view or index references tables and views that are in the grantee's own schema.
GLOBAL QUERY REWRITE	Enable rewrite using a materialized view, or create a function-based index, when that materialized view or index references tables or views in any schema.
<b>LIBRARIES</b>	
CREATE LIBRARY	Create external procedure/function libraries in grantee's schema
CREATE ANY LIBRARY	Create external procedure/function libraries in any schema
DROP LIBRARY	Drop external procedure/function libraries in the grantee's schema
DROP ANY LIBRARY	Drop external procedure/function libraries in any schema
<b>MATERIALIZED VIEWS (which are identical to SNAPSHOTS)</b>	
CREATE MATERIALIZED VIEW	Create a materialized view in the grantee's schema
CREATE ANY MATERIALIZED VIEW	Create materialized views in any schema
ALTER ANY MATERIALIZED VIEW	Alter materialized views in any schema
DROP ANY MATERIALIZED VIEW	Drop materialized views in any schema
QUERY REWRITE	Enable rewrite using a materialized view, or create a function-based index, when that materialized view or index references tables and views that are in the grantee's own schema.

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter O7\_DICTIONARY\_ACCESSIBILITY to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7–5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
GLOBAL QUERY REWRITE	Enable rewrite using a materialized view, or create a function-based index, when that materialized view or index references tables or views in any schema.
<b>OPERATORS</b>	
CREATE OPERATOR	Create an operator and its bindings in the grantee's schema
CREATE ANY OPERATOR	Create an operator and its bindings in any schema
DROP ANY OPERATOR	Drop an operator in any schema
EXECUTE ANY OPERATOR	Reference an operator in any schema
<b>OUTLINES</b>	
CREATE ANY OUTLINE	Create outlines that can be used in any schema that uses outlines
ALTER ANY OUTLINE	Modify outlines.
DROP ANY OUTLINE	Drop outlines
<b>PROCEDURES</b>	
CREATE PROCEDURE	Create stored procedures, functions, and packages in grantee's schema
CREATE ANY PROCEDURE	Create stored procedures, functions, and packages in any schema
ALTER ANY PROCEDURE	Alter stored procedures, functions, or packages in any schema
DROP ANY PROCEDURE	Drop stored procedures, functions, or packages in any schema
EXECUTE ANY PROCEDURE	Execute procedures or functions (standalone or packaged) Reference public package variables in any schema
<b>PROFILES</b>	
CREATE PROFILE	Create profiles

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter `O7_DICTIONARY_ACCESSIBILITY` to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.



**Table 7–5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
ALTER PROFILE	Alter profiles
DROP PROFILE	Drop profiles
<b>ROLES</b>	
CREATE ROLE	Create roles
ALTER ANY ROLE	Alter any role in the database
DROP ANY ROLE	Drop roles
GRANT ANY ROLE	Grant any role in the database
<b>ROLLBACK SEGMENTS</b>	
CREATE ROLLBACK SEGMENT	Create rollback segments
ALTER ROLLBACK SEGMENT	Alter rollback segments
DROP ROLLBACK SEGMENT	Drop rollback segments
<b>SEQUENCES</b>	
CREATE SEQUENCE	Create sequences in grantee's schema
CREATE ANY SEQUENCE	Create sequences in any schema
ALTER ANY SEQUENCE	Alter any sequence in the database
DROP ANY SEQUENCE	Drop sequences in any schema
SELECT ANY SEQUENCE	Reference sequences in any schema
<b>SESSIONS</b>	
CREATE SESSION	Connect to the database
ALTER RESOURCE COST	Set costs for session resources
ALTER SESSION	Issue ALTER SESSION statements
RESTRICTED SESSION	Logon after the instance is started using the SQL*Plus STARTUP RESTRICT statement

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter O7\_DICTIONARY\_ACCESSIBILITY to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7-5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
<b>SNAPSHOTS (which are identical to MATERIALIZED VIEWS)</b>	
CREATE SNAPSHOT	Create snapshots in grantee's schema
CREATE ANY SNAPSHOT	Create snapshots in any schema
ALTER ANY SNAPSHOT	Alter any snapshot in the database
DROP ANY SNAPSHOT	Drop snapshots in any schema
GLOBAL QUERY REWRITE	Enable rewrite using a snapshot, or create a function-based index, when that snapshot or index references tables or views in any schema.
QUERY REWRITE	Enable rewrite using a snapshot, or create a function-based index, when that snapshot or index references tables and views that are in the grantee's own schema.
<b>SYNONYMS</b>	
CREATE SYNONYM	Create synonyms in grantee's schema
CREATE ANY SYNONYM	Create private synonyms in any schema
CREATE PUBLIC SYNONYM	Create public synonyms
DROP ANY SYNONYM	Drop private synonyms in any schema
DROP PUBLIC SYNONYM	Drop public synonyms
<b>TABLES</b>	
CREATE ANY TABLE	Create tables in any schema. The owner of the schema containing the table must have space quota on the tablespace to contain the table.
ALTER ANY TABLE	Alter any table or view in the schema
BACKUP ANY TABLE	Use the Export utility to incrementally export objects from the schema of other users
DELETE ANY TABLE	Delete rows from tables, table partitions, or views in any schema

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter `O7_DICTIONARY_ACCESSIBILITY` to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7-5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
DROP ANY TABLE	Drop or truncate tables or table partitions in any schema
INSERT ANY TABLE	Insert rows into tables and views in any schema
LOCK ANY TABLE	Lock tables and views in any schema
UPDATE ANY TABLE	Update rows in tables and views in any schema
SELECT ANY TABLE	Query tables, views, or snapshots in any schema
<b>TABLESPACES</b>	
CREATE TABLESPACE	Create tablespaces
ALTER TABLESPACE	Alter tablespaces
DROP TABLESPACE	Drop tablespaces
MANAGE TABLESPACE	Take tablespaces offline and online and begin and end tablespace backups
UNLIMITED TABLESPACE	Use an unlimited amount of any tablespace. This privilege overrides any specific quotas assigned. If you revoke this privilege from a user, the user's schema objects remain but further tablespace allocation is denied unless authorized by specific tablespace quotas. You cannot grant this system privilege to roles.
<b>TRIGGERS</b>	
CREATE TRIGGER	Create a database trigger in grantee's schema
CREATE ANY TRIGGER	Create database triggers in any schema
ALTER ANY TRIGGER	Enable, disable, or compile database triggers in any schema
DROP ANY TRIGGER	Drop database triggers in any schema
ADMINISTER DATABASE TRIGGER	Create a trigger on DATABASE. (You must also have the CREATE TRIGGER or CREATE ANY TRIGGER privilege.)

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter O7\_DICTIONARY\_ACCESSIBILITY to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7-5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
<b>TYPES</b>	
CREATE TYPE	Create object types and object type bodies in grantee's schema
CREATE ANY TYPE	Create object types and object type bodies in any schema
ALTER ANY TYPE	Alter object types in any schema
DROP ANY TYPE	Drop object types and object type bodies in any schema
EXECUTE ANY TYPE	Use and reference object types and collection types in any schema, and invoke methods of an object type in any schema <i>if you make the grant to a specific user</i> . If you grant EXECUTE ANY TYPE to a role, users holding the enabled role will not be able to invoke methods of an object type in any schema.
<b>USERS</b>	
CREATE USER	Create users. This privilege also allows the creator to <ul style="list-style-type: none"> <li>■ Assign quotas on <b>any</b> tablespace,</li> <li>■ Set default and temporary tablespaces, and</li> <li>■ Assign a profile as part of a CREATE USER statement.</li> </ul>
ALTER USER	Alter any user. This privilege authorizes the grantee to <ul style="list-style-type: none"> <li>■ Change another user's password or authentication method,</li> <li>■ Assign quotas on <b>any</b> tablespace,</li> <li>■ Set default and temporary tablespaces, and</li> <li>■ Assign a profile and default roles</li> </ul>
BECOME USER	Become another user. (Required by any user performing a full database import.)
DROP USER	Drop users

**Note:** When you grant a privilege on "ANY" object (for example, CREATE ANY CLUSTER), you give the user access to that type of object in all schemas, including the SYS schema. If you want to prohibit access to objects in the SYS schema, set the initialization parameter O7\_DICTIONARY\_ACCESSIBILITY to FALSE. Then privileges granted on "ANY" object will allow access to any schema except SYS.

**Table 7-5 (Cont.) System Privileges**

<b>System Privilege</b>	<b>Allows grantee to . . .</b>
<b>VIEWS</b>	
CREATE VIEW	Create views in grantee's schema
CREATE ANY VIEW	Create views in any schema
DROP ANY VIEW	Drop views in any schema
<b>MISCELLANEOUS</b>	
ANALYZE ANY	Analyze any table, cluster, or index in any schema
AUDIT ANY	Audit any object in any schema using <code>AUDIT <i>schema_</i>objects</code> statements
COMMENT ANY TABLE	Comment on any table, view, or column in any schema
FORCE ANY TRANSACTION	Force the commit or rollback of any in-doubt distributed transaction in the local database Induce the failure of a distributed transaction
FORCE TRANSACTION	Force the commit or rollback of grantee's in-doubt distributed transactions in the local database
GRANT ANY PRIVILEGE	Grant any system privilege.
SYSDBA	Perform <code>STARTUP</code> and <code>SHUTDOWN</code> operations <code>ALTER DATABASE: open, mount, back up, or change character set</code> <code>CREATE DATABASE</code> <code>ARCHIVELOG</code> and <code>RECOVERY</code> Includes the <code>RESTRICTED SESSION</code> privilege
SYSOPER	Perform <code>STARTUP</code> and <code>SHUTDOWN</code> operations <code>ALTER DATABASE OPEN/MOUNT/BACKUP</code> <code>ARCHIVELOG</code> and <code>RECOVERY</code> Includes the <code>RESTRICTED SESSION</code> privilege

**Note:** When you grant a privilege on "ANY" object (for example, `CREATE ANY CLUSTER`), you give the user access to that type of object in all schemas, including the `SYS` schema. If you want to prohibit access to objects in the `SYS` schema, set the initialization parameter `O7_DICTIONARY_ACCESSIBILITY` to `FALSE`. Then privileges granted on "ANY" object will allow access to any schema except `SYS`.

**Table 7–6 Oracle Predefined Roles**

Predefined Role	Purpose
CONNECT, RESOURCE, and DBA	<p>These roles are provided for compatibility with previous versions of Oracle. You can determine the privileges encompassed by these roles by querying the DBA_SYS_PRIVILEGES data dictionary view. See <i>Oracle8i Reference</i> for a description of this view.</p> <p><b>Note:</b> Oracle Corporation recommends that you design your own roles for database security rather than relying on these roles. These roles may not be created automatically by future versions of Oracle.</p>
DELETE_CATALOG_ROLE EXECUTE_CATALOG_ROLE SELECT_CATALOG_ROLE	<p>These roles are provided for accessing data dictionary views and packages.</p> <p><b>See Also:</b> <i>Oracle8i Administrator's Guide</i> for more information on these roles.</p>
EXP_FULL_DATABASE IMP_FULL_DATABASE	<p>These roles are provided for convenience in using the Import and Export utilities.</p> <p><b>See Also:</b> <i>Oracle8i Utilities</i> for more information on these roles.</p>
AQ_USER_ROLE AQ_ADMINISTRATOR_ROLE	<p>You need these roles to use Oracle's Advanced Queuing functionality.</p> <p><b>See Also:</b> <i>Oracle8i Application Developer's Guide - Advanced Queuing</i> for more information on these roles.</p>
SNMPAGENT	<p>This role is used by Enterprise Manager/Intelligent Agent.</p> <p><b>See Also:</b> <i>Oracle Enterprise Manager Administrator's Guide</i>.</p>
RECOVERY_CATALOG_ OWNER	<p>You need this role to create a user who owns a recovery catalog.</p> <p><b>See Also:</b> <i>Oracle8i Backup and Recovery Guide</i> for more information on recovery catalogs.</p>
HS_ADMIN_ROLE	<p>A DBA using Oracle's heterogeneous services feature needs this role to access appropriate tables in the data dictionary and to manipulate them with the DBMS_HS package.</p> <p><b>See Also:</b> <i>Oracle8i Distributed Database Systems</i> and <i>Oracle8i Supplied PL/SQL Packages Reference</i> for more information.</p>
<p>Oracle also creates other roles that authorize you to administer the database. On many operating systems, these roles are called OSOPER and OSDBA. Their names may be different on your operating system.</p>	

## Examples

To grant the `CREATE SESSION` system privilege to `RICHARD`, allowing `RICHARD` to log on to Oracle, issue the following statement:

```
GRANT CREATE SESSION
  TO richard;
```

To grant the `CREATE TABLE` system privilege to the role `TRAVEL_AGENT`, issue the following statement:

```
GRANT CREATE TABLE
  TO travel_agent;
```

`TRAVEL_AGENT`'s privilege domain now contains the `CREATE TABLE` system privilege.

The following statement grants the `TRAVEL_AGENT` role to the `EXECUTIVE` role:

```
GRANT travel_agent
  TO executive;
```

`TRAVEL_AGENT` is now granted to `EXECUTIVE`. `EXECUTIVE`'s privilege domain contains the `CREATE TABLE` system privilege.

To grant the `EXECUTIVE` role with the `ADMIN OPTION` to `THOMAS`, issue the following statement:

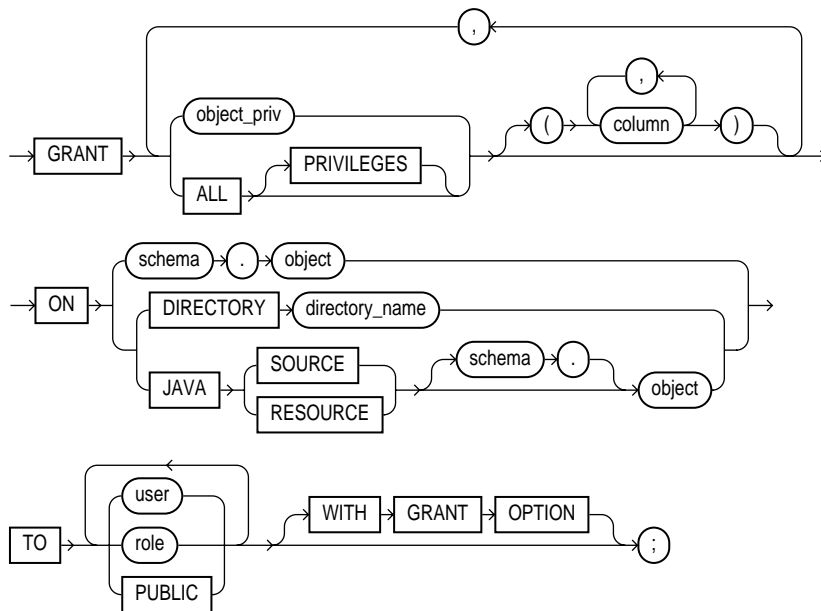
```
GRANT executive
  TO thomas
  WITH ADMIN OPTION;
```

`THOMAS` can now perform the following operations with the `EXECUTIVE` role:

- Enable the role and exercise any privileges in the role's privilege domain, including the `CREATE TABLE` system privilege
- Grant and revoke the role to and from other users
- Drop the role

## GRANT *object\_privileges*

### Syntax



### Purpose

To grant privileges for a particular object to users, roles, and `PUBLIC`. To grant system privileges and roles, use the `GRANT system_privileges_and_roles` statement described in the previous section of this chapter. [Table 7-7](#) summarizes the object privileges that you can grant on each type of object.

If you grant a **privilege to a user**, Oracle adds the privilege to the user's privilege domain. The user can immediately exercise the privilege.

If you grant a **privilege to a role**, Oracle adds the privilege to the role's privilege domain. Users who have been granted and have enabled the role can immediately exercise the privilege. Other users who have been granted the role can enable the role and exercise the privilege.



If you grant a **privilege to PUBLIC**, Oracle adds the privilege to the privilege domain of each user. All users can immediately exercise the privilege.

[Table 7-8](#) lists object privileges and the operations that they authorize. You can grant any of these system privileges with the `GRANT` statement.

For information on granting system privileges and roles, see "[GRANT system\\_privileges\\_and\\_roles](#)" on page 7-519. For information on revoking object grants, see "[REVOKE schema\\_object\\_privileges](#)" on page 7-559.

## Prerequisites

You must own the object or the owner of the object must have granted you the object privileges with the `GRANT OPTION`. This rule applies to users with the `DBA` role.

## Keywords and Parameters

---

<i>object_priv</i>	is an object privilege to be granted. You can substitute any of the values shown in <a href="#">Table 7-7</a> . See also <a href="#">Table 7-8</a> .
	<b>Restriction:</b> A privilege cannot appear more than once in the list of privileges to be granted.
ALL [ PRIVILEGES ]	grants all the privileges for the object that you have been granted with the <code>GRANT OPTION</code> . The user who owns the schema containing an object automatically has all privileges on the object with the <code>GRANT OPTION</code> . (The keyword <code>PRIVILEGES</code> is optional.)
<i>column</i>	specifies a table or view column on which privileges are granted. You can specify columns only when granting the <code>INSERT</code> , <code>REFERENCES</code> , or <code>UPDATE</code> privilege. If you do not list columns, the grantee has the specified privilege on all columns in the table or view.
ON	identifies the object on which the privileges are granted. Directory schema objects and Java source and resource schema objects are identified separately because they reside in separate namespaces.
<i>object</i>	identifies the schema object on which the privileges are granted. If you do not qualify <i>object</i> with <i>schema</i> , Oracle assumes the object is in your own schema. The object can be one of the following types (see <a href="#">Table 7-7</a> ): <ul style="list-style-type: none"> <li>■ table, view, or materialized view / snapshot</li> <li>■ sequence</li> <li>■ procedure, function, or package</li> <li>■ user-defined type</li> <li>■ synonym for any of the above items</li> <li>■ directory, library, operator, or indextype</li> <li>■ a Java source, class, or resource</li> </ul>

---

	<b>Note:</b> You cannot grant privileges directly to a single partition of a partitioned table. For information on how to grant privileges to a single partition indirectly, refer to <i>Oracle8i Concepts</i> .
DIRECTORY	identifies a directory schema object on which privileges are granted by the DBA. You cannot qualify <i>directory_name</i> with a schema name. See " <a href="#">CREATE DIRECTORY</a> " on page 7-282.
JAVA_SOURCE   RESOURCE	identifies a Java source or resource schema object on which privileges are granted. See " <a href="#">CREATE JAVA</a> " on page 7-311.
TO	identifies users or roles to which the object privilege is granted. <b>Restriction:</b> A user or role cannot appear more than once in the TO clause.
PUBLIC	grants object privileges to all users.
WITH GRANT OPTION	allows the grantee to grant the object privileges to other users and roles. <b>Restriction:</b> You can specify this clause only when granting to a user or to PUBLIC, not when granting to a role.

---

**Table 7-7 Object Privileges**

Object Privilege	Table	View	Sequence	Procedures, Functions, Packages <sup>a</sup>	Materialized View / Snapshot	Directory	Library	User-defined Type	Operator	Index-type
ALTER	X		X							
DELETE	X	X			X <sup>b</sup>					
EXECUTE				X			X	X	X	X
INDEX	X									
INSERT	X	X			X <sup>b</sup>					
READ						X				
REFERENCES	X									
SELECT	X	X	X		X					
UPDATE	X	X			X <sup>b</sup>					

<sup>a</sup>Oracle treats a Java class, source, or resource as if it were a procedure for purposes of granting object privileges.

<sup>b</sup>The DELETE, INSERT, and UPDATE privileges can be granted only to *updatable* materialized views.

**Table 7-8 Object Privileges and the Operations They Authorize**

Object Privilege	Allows Grantee to . . .
The following <b>table privileges</b> authorize operations on a table. Any one of following object privileges allows the grantee to lock the table in any lock mode with the LOCK TABLE statement.	
ALTER	Change the table definition with the ALTER TABLE statement.
DELETE	Remove rows from the table with the DELETE statement. <b>Note:</b> You must grant the SELECT privilege on the table along with the DELETE privilege.
INDEX	Create an index on the table with the CREATE INDEX statement.
INSERT	Add new rows to the table with the INSERT statement.
REFERENCES	Create a constraint that refers to the table. You cannot grant this privilege to a role.
SELECT	Query the table with the SELECT statement.

**Table 7–8 (Cont.) Object Privileges and the Operations They Authorize**

<b>Object Privilege</b>	<b>Allows Grantee to . . .</b>
UPDATE	Change data in the table with the UPDATE statement.
	<b>Note:</b> You must grant the SELECT privilege on the table along with the UPDATE privilege.

The following **view privileges** authorize operations on a view. Any one of the following object privileges allows the grantee to lock the view in any lock mode with the LOCK TABLE statement.

To grant a privilege on a view, you must have that privilege with the GRANT OPTION on all of the view's base tables.

DELETE	Remove rows from the view with the DELETE statement.
INSERT	Add new rows to the view with the INSERT statement.
SELECT	Query the view with the SELECT statement.
UPDATE	Change data in the view with the UPDATE statement.

The following **sequence privileges** authorize operations on a sequence.

ALTER	Change the sequence definition with the ALTER SEQUENCE statement.
SELECT	Examine and increment values of the sequence with the CURRVAL and NEXTVAL pseudocolumns.

The following **procedure, function, and package privilege** authorizes operations on procedures, functions, or packages. This privilege also applies to **Java sources, classes, and resources**, which Oracle treats as though they were procedures for purposes of granting object privileges.

EXECUTE	Compile the procedure or function or execute it directly, or access any program object declared in the specification of a package.
	<b>Note:</b> Users do not need this privilege to execute a procedure, function, or package indirectly.
	<b>See Also:</b> <i>Oracle8i Concepts</i> and <i>Oracle8i Application Developer's Guide - Fundamentals</i> .

The following **snapshot privilege** authorizes operations on a snapshot.

SELECT	Query the snapshot with the SELECT statement.
--------	---

**Table 7–8 (Cont.) Object Privileges and the Operations They Authorize**

Object Privilege	Allows Grantee to . . .
<p><b>Synonym privileges</b> are the same as the privileges for the base object. Granting a privilege on a synonym is equivalent to granting the privilege on the base object. Similarly, granting a privilege on a base object is equivalent to granting the privilege on all synonyms for the object. If you grant a user a privilege on a synonym, the user can use either the synonym name or the base object name in the SQL statement that exercises the privilege.</p> <p>The following <b>directory privilege</b> provides secured access to the files stored in the operating system directory to which the directory object serves as a pointer. The directory object contains the full pathname of the operating system directory where the files reside. Because the files are actually stored outside the database, Oracle server processes also need to have appropriate file permissions on the file system server. Granting object privileges on the directory database object to individual database users, rather than on the operating system, allows Oracle to enforce security during file operations.</p>	
READ	Read files in the directory.
<p>The following <b>object type privilege</b> authorizes operations on an object type</p>	
EXECUTE	Use and reference the specified object and to invoke its methods.
<p>The following <b>indextype privilege</b> authorizes operations on indextypes.</p>	
EXECUTE	Reference an indextype.
<p>The following <b>operator privilege</b> authorizes operations on user-defined operators.</p>	
EXECUTE	Reference an operator.

## Examples

To grant READ on directory BFILE\_DIR1 to user SCOTT, with the GRANT OPTION, issue the following statement:

```
GRANT READ ON DIRECTORY bfile_dir1 TO scott
  WITH GRANT OPTION;
```

To grant all privileges on the table BONUS to the user JONES with the GRANT OPTION, issue the following statement:

```
GRANT ALL ON bonus TO jones
  WITH GRANT OPTION;
```

JONES can subsequently perform the following operations:

- Exercise any privilege on the BONUS table
- Grant any privilege on the BONUS table to another user or role

To grant **SELECT** and **UPDATE** privileges on the view **GOLF\_HANDICAP** to all users, issue the following statement:

```
GRANT SELECT, UPDATE
  ON golf_handicap TO PUBLIC;
```

All users can subsequently query and update the view of golf handicaps.

To grant **SELECT** privilege on the **ESEQ** sequence in the schema **ELLY** to the user **BLAKE**, issue the following statement:

```
GRANT SELECT
  ON elly.eseq TO blake;
```

**BLAKE** can subsequently generate the next value of the sequence with the following statement:

```
SELECT elly.eseq.NEXTVAL
  FROM DUAL;
```

To grant **BLAKE** the **REFERENCES** privilege on the **EMPNO** column and the **UPDATE** privilege on the **EMPNO**, **SAL**, and **COMM** columns of the **EMP** table in the schema **SCOTT**, issue the following statement:

```
GRANT REFERENCES (empno), UPDATE (empno, sal, comm)
  ON scott.emp
  TO blake;
```

**BLAKE** can subsequently update values of the **EMPNO**, **SAL**, and **COMM** columns. **BLAKE** can also define referential integrity constraints that refer to the **EMPNO** column. However, because the **GRANT** statement lists only these columns, **BLAKE** cannot perform operations on any of the other columns of the **EMP** table.


For example, **BLAKE** can create a table with a constraint:

```
CREATE TABLE dependent
  (dependno    NUMBER,
   dependname  VARCHAR2(10),
   employee    NUMBER
   CONSTRAINT in_emp REFERENCES scott.emp(empno) );
```

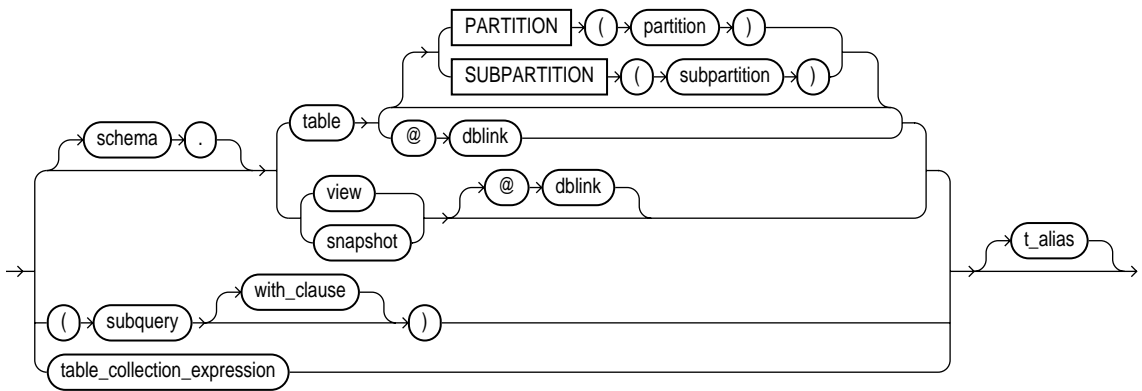
The constraint **IN\_EMP** ensures that all dependents in the **DEPENDENT** table correspond to an employee in the **EMP** table in the schema **SCOTT**.

# INSERT

## Syntax




### DML\_table\_expression\_clause::=




subquery: see "SELECT and Subqueries" on page 7-569.

### with\_clause::=




### table\_collection\_expression::=




## INSERT

---

### values\_clause::=



### returning\_clause::=



## Purpose

To add rows to a table, a view's base table, a partition of a partitioned table or a subpartition of a composite-partitioned table, or an object table or an object view's base table.

## Prerequisites

For you to insert rows into a table, the table must be in your own schema or you must have `INSERT` privilege on the table.

For you to insert rows into the base table of a view, the owner of the schema containing the view must have `INSERT` privilege on the base table. Also, if the view is in a schema other than your own, you must have `INSERT` privilege on the view.

If you have the `INSERT ANY TABLE` system privilege, you can also insert rows into any table or any view's base table.

## Keywords and Parameters

---

*hint* is a comment that passes instructions to the optimizer on choosing an execution plan for the statement. For the syntax and description of hints, see "[Hints](#)" on page 2-63 and *Oracle8i Designing and Tuning for Performance*.

*DML\_table\_expression\_clause*

*schema* is the schema containing the table or view. If you omit *schema*, Oracle assumes the table or view is in your own schema.



*table* | *view* |  
*subquery*

is the name of the table or object table, or view or object view, or the column or columns returned by a subquery, into which rows are to be inserted. If you specify a view or object view, Oracle inserts rows into the view's base table.

If any value to be inserted is a REF to an object table, and if the object table has a primary key object identifier, then the column into which you insert the REF must be a REF column with a referential integrity or SCOPE constraint to the object table.

If *table* (or the base table of *view*) contains one or more domain index columns, this statement executes the appropriate indextype insert routine.

Issuing an INSERT statement against a table fires any INSERT triggers defined on the table.

**See Also:** *Oracle8i Data Cartridge Developer's Guide* for more information on these routines

#### Restrictions:

- You cannot execute this statement if *table* (or the base table of *view*) contains any domain indexes marked LOADING or FAILED.
- With regard to the ORDER BY clause of the *subquery* in the *DML\_query\_expression\_clause*, ordering is guaranteed only for the rows being inserted, and only within each extent of the table. Ordering of new rows with respect to existing rows is not guaranteed.
- If a view was created using the WITH CHECK OPTION, then you can insert into the view only rows that satisfy the view's defining query.
- If a view was created using a single base table, then you can insert rows into the view and then retrieve those values using the *returning\_clause*.
- You cannot insert rows into a view except with INSTEAD OF triggers if the view's defining query contains one of the following constructs:
  - A set operator
  - A DISTINCT operator
  - An aggregate function
  - A GROUP BY, ORDER BY, CONNECT BY, or START WITH clause
  - A collection expression in a SELECT list
  - A subquery in a SELECT list
  - Joins (with some exceptions). See *Oracle8i Administrator's Guide* for details.
- If you specify an index, index partition, or index subpartition that has been marked UNUSABLE, the INSERT statement will fail unless the SKIP\_UNUSABLE\_INDEXES parameter has been set to TRUE.

**See Also:** "ALTER SESSION" on page 7-83.

---

<code>PARTITION</code> ( <i>partition_name</i> ) / <code>SUBPARTITION</code> ( <i>subpartition_name</i> )	<p>specifies the name of the partition or subpartition within <i>table</i> (or the base table of <i>view</i>) targeted for inserts.</p> <p>If a row to be inserted does not map into a specified partition or subpartition, Oracle returns an error.</p> <p><b>Restriction:</b> This clause is not valid for object tables or object views.</p>
<i>dblink</i>	<p>is a complete or partial name of a database link to a remote database where the table or view is located. For information on referring to database links, see "<a href="#">Referring to Schema Objects and Parts</a>" on page 2-76. You can insert rows into a remote table or view only if you are using Oracle's distributed functionality.</p> <p>If you omit <i>dblink</i>, Oracle assumes that the table or view is on the local database.</p>
<i>with_clause</i>	<p>restricts the subquery in one of the following ways:</p> <ul style="list-style-type: none"><li>▪ <code>WITH READ ONLY</code> specifies that the subquery cannot be updated.</li><li>▪ <code>WITH CHECK OPTION</code> specifies that Oracle prohibits any changes to that table that would produce rows that are not included in the subquery. See the "<a href="#">WITH CHECK OPTION Example</a>" on page 7-587.</li></ul>
<i>table_collection_expression</i>	<p>informs Oracle that the collection value expression should be treated as a table. See "<a href="#">Table Collection Examples</a>" on page 7-586.</p> <p><i>collection_expression</i> is a subquery that selects a nested table column from <i>table</i> or <i>view</i>.</p> <hr/> <p><b>Note:</b> In earlier releases of Oracle, <i>table_collection_expression</i> was expressed as "THE subquery". That usage is now deprecated.</p> <hr/>
<i>column</i>	<p>is a column of the table or view. In the inserted row, each column in this list is assigned a value from the <i>values_clause</i> or the subquery.</p> <p>If you omit one of the table's columns from this list, the column's value for the inserted row is the column's default value as specified when the table was created. If any of these columns has a <code>NOT NULL</code> constraint, then Oracle returns an error indicating that the constraint has been violated and rolls back the <code>INSERT</code> statement.</p> <p>If you omit the column list altogether, the <i>values_clause</i> or query must specify values for all columns in the table.</p> <p><b>See Also:</b> "<a href="#">CREATE TABLE</a>" on page 4-381 for more information on default column values.</p>
<i>values_clause</i>	<p>specifies a row of values to be inserted into the table or view. See the syntax description in "<a href="#">Expressions</a>" on page 5-1 and "<a href="#">SELECT and Subqueries</a>" on page 7-569. You must specify a value in the <i>values_clause</i> for each column in the column list. If you omit the column list, then the <i>values_clause</i> must provide values for every column in the table.</p>

**Restrictions:**

- You cannot initialize an internal LOB attribute in an object with a value other than empty or null. That is, you cannot use a literal.
- You cannot insert a BFILE value until you have initialized the BFILE locator to null or to a directory alias and filename. See the "[BFILE Example](#)" on page 7-546. For information on initializing BFILES, see *Oracle Call Interface Programmer's Guide* and *Oracle8i Application Developer's Guide - Fundamentals*.

**Note:** If you insert string literals into a RAW column, during subsequent queries, Oracle will perform a full table scan rather than using any index that might exist on the RAW column.

*subquery*

is a subquery that returns rows that are inserted into the table. If the subquery selects no rows, Oracle inserts no rows into the table.

- When specified with VALUES, the subquery returns values to be inserted into one row.
- When specified without VALUES, the subquery can return values to be inserted into more than one row.

The subquery can refer to any table, view, or snapshot, including the target table of the INSERT statement. The select list of this subquery must have the same number of columns as the column list of the INSERT statement. If you omit the column list, then the subquery must provide values for every column in the table. See "[SELECT and Subqueries](#)" on page 7-569.

You can use *subquery* in combination with the TO\_LOB function to convert the values in a LONG column to LOB values in another column in the same or another table. For a discussion of why and when to copy LONGs to LOBs, see *Oracle8i Migration*. For a description of how to use the TO\_LOB function, see "[Conversion Functions](#)" on page 4-5. See also the [TO\\_LOB Example](#) on page 7-546. To migrate LONGs to LOBs in a view, you must perform the migration on the base table, and then add the LOB to the view.

**Notes:**

- If *subquery* returns (in part or totally) the equivalent of an existing materialized view, Oracle may use the materialized view (for query rewrite) in place of one or more tables specified in *subquery*.  
**See Also:** *Oracle8i Data Warehousing Guide* for more information on materialized views and query rewrite.
- If this *subquery* refers to remote objects, the INSERT operation can run in parallel as long as the reference does not loop back to an object on the local database. (See the description of the *parallel\_clause* in "[CREATE TABLE](#)" on page 4-381.) However, if the *subquery* in the *DML\_query\_expression\_clause* refers to any remote objects, the INSERT operation will run serially without notification.

*t\_alias*

provides a **correlation name** for the table, view, or subquery to be referenced elsewhere in the statement.

*returning\_clause* retrieves the rows affected by the INSERT.

An INSERT statement with a *returning\_clause* retrieves the rows inserted and stores them in PL/SQL variables or bind variables. Using a *returning\_clause* in INSERT statements with a *values\_clause* enables you to return column expressions, ROWIDs, and REFS and store them in output bind variables. You can also use INSERT with a *returning\_clause* for views with single base tables.

*expr* is some form of the syntax descriptions in "Expressions" on page 5-1. You must specify a column expression in the *returning\_clause* for each variable in the *data\_item\_list*.

INTO indicates that the values of the changed rows are to be stored in the variable(s) specified in *data\_item\_list*.

*data\_item* is a PL/SQL variable or bind variable that stores a retrieved *expr* value.

**Restrictions:**

- You cannot use this clause with parallel DML or with remote objects.
  - You cannot retrieve LONG types with this clause.
  - You cannot specify this clause for a view on which an INSTEAD OF trigger has been defined.
- 

## Examples

**VALUES Examples** The following statement inserts a row into the DEPT table:

```
INSERT INTO dept
VALUES (50, 'PRODUCTION', 'SAN FRANCISCO');
```

The following statement inserts a row with six columns into the EMP table. One of these columns is assigned NULL and another is assigned a number in scientific notation:

```
INSERT INTO emp (empno, ename, job, sal, comm, deptno)
VALUES (7890, 'JINKS', 'CLERK', 1.2E3, NULL, 40);
```

The following statement has the same effect as the preceding example, but uses a subquery in the *DML\_query\_expression\_clause*:

```
INSERT INTO (SELECT empno, ename, job, sal, comm, deptno FROM emp)
VALUES (7890, 'JINKS', 'CLERK', 1.2E3, NULL, 40);
```

**Subquery Example** The following statement copies managers and presidents or employees whose commission exceeds 25% of their salary into the BONUS table:

```
INSERT INTO bonus
```

```

SELECT ename, job, sal, comm
FROM emp
WHERE comm > 0.25 * sal
OR job IN ('PRESIDENT', 'MANAGER');

```

**Database Link Example** The following statement inserts a row into the ACCOUNTS table owned by the user SCOTT on the database accessible by the database link SALES:

```

INSERT INTO scott.accounts@sales (acc_no, acc_name)
VALUES (5001, 'BOWER');

```

Assuming that the ACCOUNTS table has a BALANCE column, the newly inserted row is assigned the default value for this column (if one has been defined), because this INSERT statement does not specify a BALANCE value.

**Sequence Example** The following statement inserts a new row containing the next value of the employee sequence into the EMP table:

```

INSERT INTO emp
VALUES (empseq.nextval, 'LEWIS', 'CLERK',
       7902, SYSDATE, 1200, NULL, 20);

```

**Partition Example** The following example adds rows from LATEST\_DATA into partition OCT98 of the SALES table:

```

INSERT INTO sales PARTITION (oct98)
SELECT * FROM latest_data;

```

**Bind Variable Example** The following example returns the values of the inserted rows into output bind variables :BND1 and :BND2:

```

INSERT INTO emp VALUES (empseq.nextval, 'LEWIS', 'CLARK',
                       7902, SYSDATE, 1200, NULL, 20)
RETURNING sal*12, job INTO :bnd1, :bnd2;

```

**Bind Array Example** The following example returns the reference value for the inserted row into bind array :1:

```

INSERT INTO employee
VALUES ('Kitty Mine', 'Peaches Fuzz', 'Meena Katz')
RETURNING REF(employee) INTO :1;

```

**TO\_LOB Example** The following example copies LONG data to a LOB column in the following existing table:

```
CREATE TABLE long_tab (long_pics LONG RAW);
```

First you must create a table with a LOB.

```
CREATE TABLE lob_tab (lob_pics BLOB);
```

Next, use an INSERT ... SELECT statement to copy the data in all rows for the LONG column into the newly created LOB column:

```
INSERT INTO lob_tab (lob_pics)
  SELECT TO_LOB(long_pics) FROM long_tab;
```

Once you are confident that the migration has been successful, you can drop the LONG\_PICS table. Alternatively, if the table contains other columns, you can simply drop the LONG column from the table as follows:

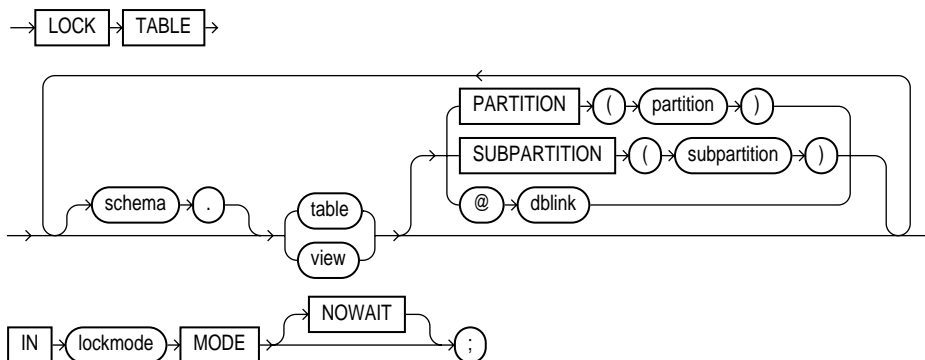
```
ALTER TABLE long_tab DROP COLUMN long_pics;
```

**BFILE Example** When you INSERT or UPDATE a BFILE, you must initialize it to null or to a directory alias and filename, as shown in the next example. Assume that the EMP table has a NUMBER column followed by a BFILE column:

```
INSERT INTO emp
  VALUES (1, BFILENAME ('a_dir_alias', 'a_filename'));
```

## LOCK TABLE

### Syntax



### Purpose

To lock one or more tables (or table partitions or subpartitions) in a specified mode. This lock manually overrides automatic locking and permits or denies access to a table or view by other users for the duration of your operation.

Some forms of locks can be placed on the same table at the same time. Other locks allow only one lock per table. For a complete description of the interaction of lock modes, see *Oracle8i Concepts*.

A locked table remains locked until you either commit your transaction or roll it back, either entirely or to a savepoint before you locked the table.

A lock never prevents other users from querying the table. A query never places a lock on a table. Readers never block writers and writers never block readers.

**See Also:** ["COMMIT"](#) on page 7-230, ["ROLLBACK"](#) on page 7-564, and ["SAVEPOINT"](#) on page 7-567.

### Prerequisites

The table or view must be in your own schema or you must have the LOCK ANY TABLE system privilege, or you must have any object privilege on the table or view.

## Keywords and Parameters

<i>schema</i>	is the schema containing the table or view. If you omit <i>schema</i> , Oracle assumes the table or view is in your own schema.
<i>table / view</i>	<p>is the name of the table to be locked. If you specify <i>view</i>, Oracle locks the view's base tables.</p> <p>If you specify <code>PARTITION</code> (<i>partition</i>) or <code>SUBPARTITION</code> (<i>subpartition</i>), Oracle first acquires an implicit lock on the table. The table lock is the same as the lock you specify for <i>partition</i> or <i>subpartition</i>, with two exceptions:</p> <ul style="list-style-type: none"><li>■ If you specify a <code>SHARE</code> lock for the subpartition, Oracle acquires an implicit <code>ROW SHARE</code> lock on the table.</li><li>■ If you specify an <code>EXCLUSIVE</code> lock for the subpartition, Oracle acquires an implicit <code>ROW EXCLUSIVE</code> lock on the table.</li></ul> <p>If you specify <code>PARTITION</code> and <i>table</i> is composite-partitioned, then Oracle acquires locks on all the subpartitions of <i>partition</i>.</p>
<i>dblink</i>	<p>is a database link to a remote Oracle database where the table or view is located. For information on specifying database links, see <a href="#">"Referring to Objects in Remote Databases"</a> on page 2-79. You can lock tables and views on a remote database only if you are using Oracle's distributed functionality. All tables locked by a <code>LOCK TABLE</code> statement must be on the same database.</p> <p>If you omit <i>dblink</i>, Oracle assumes the table or view is on the local database.</p>
<i>lockmode</i>	<p>is one of the following:</p> <p><code>ROW SHARE</code> allows concurrent access to the locked table, but prohibits users from locking the entire table for exclusive access. <code>ROW SHARE</code> is synonymous with <code>SHARE UPDATE</code>, which is included for compatibility with earlier versions of Oracle.</p> <p><code>ROW EXCLUSIVE</code> is the same as <code>ROW SHARE</code>, but also prohibits locking in <code>SHARE</code> mode. Row Exclusive locks are automatically obtained when updating, inserting, or deleting.</p> <p><code>SHARE UPDATE</code>—see <code>ROW SHARE</code>.</p> <p><code>SHARE</code> allows concurrent queries but prohibits updates to the locked table.</p> <p><code>SHARE ROW EXCLUSIVE</code> is used to look at a whole table and to allow others to look at rows in the table but to prohibit others from locking the table in <code>SHARE</code> mode or updating rows.</p> <p><code>EXCLUSIVE</code> allows queries on the locked table but prohibits any other activity on it.</p>
<code>NOWAIT</code>	specifies that Oracle returns control to you immediately if the specified table (or specified partition or subpartition) is already locked by another user. In this case, Oracle returns a message indicating that the table, partition, or subpartition is already locked by another user.



---

If you omit this clause, Oracle waits until the table is available, locks it, and returns control to you.

---

## Examples

The following statement locks the EMP table in exclusive mode, but does not wait if another user already has locked the table:

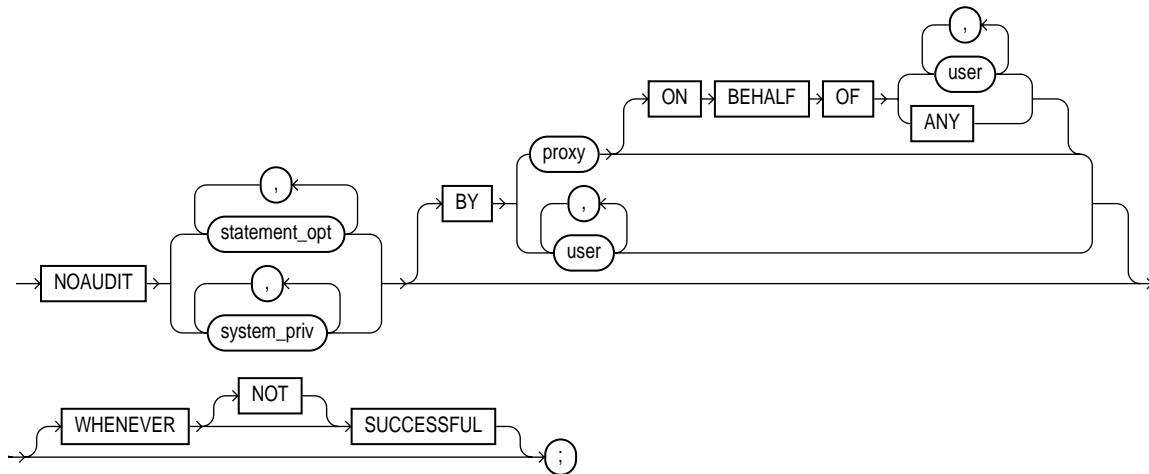
```
LOCK TABLE emp
  IN EXCLUSIVE MODE
  NOWAIT;
```

The following statement locks the remote ACCOUNTS table that is accessible through the database link BOSTON:

```
LOCK TABLE accounts@boston
  IN SHARE MODE;
```

## NOAUDIT *sql\_statements*

### Syntax



### Purpose

To stop auditing previously enabled by the `AUDIT sql_statements` statement. To stop auditing enabled by the `AUDIT schema_objects` statement, refer to "[NOAUDIT \*schema\\_objects\*](#)" on page 7-552.

The `NOAUDIT` statement must have the same syntax as the previous `AUDIT` statement. Further, it reverses the effects only of that particular statement. Therefore, if one `AUDIT` statement (statement A) enables auditing for a specific user, and a second (statement B) enables auditing for all users, then a `NOAUDIT` statement to disable auditing for all users (statement C) reverses statement B, but leaves statement A in effect and continues to audit the user that statement A specified. For information on auditing specific SQL statements, see the "[AUDIT \*sql\\_statements\*](#)" on page 7-213.

### Prerequisites

You must have the `AUDIT SYSTEM` system privilege.

## Keywords and Parameters

---

<i>statement_opt</i>	is a statement option for which auditing is stopped. For a list of the statement options and the SQL statements they audit, see <a href="#">Table 7-1</a> on page 7-215 and <a href="#">Table 7-2</a> on page 7-218.
<i>system_priv</i>	is a system privilege for which auditing is stopped. For a list of the system privileges and the statements they authorize, see <a href="#">Table 7-5</a> on page 7-521.
BY <i>user</i>	stops auditing only for SQL statements issued by specified users in their subsequent sessions. If you omit this clause, Oracle stops auditing for all users' statements, except for the situation described for <code>WHENEVER SUCCESSFUL</code> .
BY <i>proxy</i>	stops auditing only for the SQL statements issued by the specified proxy, on behalf of a specific user or any user.
WHENEVER SUCCESSFUL	stops auditing only for SQL statements that complete successfully. NOT stops auditing only for statements that result in Oracle errors.  If you omit the <code>WHENEVER SUCCESSFUL</code> clause entirely, Oracle stops auditing for all statements, regardless of success or failure.

---

## Examples

The following examples correspond to three examples listed in "[AUDIT sql\\_statements](#)" on page 7-213.

If you have chosen auditing for every SQL statement that creates or drops a role, you can stop auditing of such statements by issuing the following statement:

```
NOAUDIT ROLE;
```

If you have chosen auditing for any statement that queries or updates any table issued by the users `SCOTT` and `BLAKE`, you can stop auditing for `SCOTT`'s queries by issuing the following statement:

```
NOAUDIT SELECT TABLE
  BY scott;
```

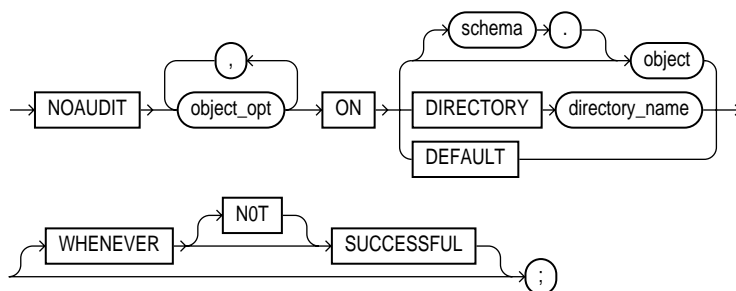
The above statement stops auditing only `SCOTT`'s queries, so Oracle continues to audit `BLAKE`'s queries and updates as well as `SCOTT`'s updates.

To stop auditing on all statements that are authorized by `DELETE ANY TABLE` system privilege, issue the following statement:

```
NOAUDIT DELETE ANY TABLE;
```

## NOAUDIT *schema\_objects*

### Syntax



### Purpose

To stop auditing previously enabled by the `AUDIT schema_objects` statement.

#### See Also:

- ["AUDIT \*schema\\_objects\*"](#) on page 7-221 for more information on auditing.
- ["NOAUDIT \*sql\\_statements\*"](#) on page 7-550 for information on stopping auditing enabled by the `AUDIT sql_statements` statement

### Prerequisites

The object on which you stop auditing must be in your own schema or you must have the `AUDIT ANY` system privilege. In addition, if the object you chose for auditing is a directory, even if you created it, you must have the `AUDIT ANY` system privilege.

### Keywords and Parameters

---

<i>object_opt</i>	stops auditing for particular operations on the object. For a list of these options, see <a href="#">Table 7-3</a> on page 7-223.
ON	identifies the object on which auditing is stopped. If you do not qualify <i>object</i> with <i>schema</i> , Oracle assumes the object is in your own schema.

---

<i>object</i>	must a table, view, sequence, stored procedure, function, or package, snapshot, or library. For information on auditing specific schema objects, refer to " <a href="#">AUDIT schema_objects</a> " on page 7-221.
DIRECTORY <i>directory_name</i>	identifies the name of the directory on which auditing is being stopped.
DEFAULT	removes the specified object options as default object options for subsequently created objects.
WHENEVER SUCCESSFUL	stops auditing only for SQL statements that complete successfully. NOT stops auditing only for statements that result in Oracle errors.  If you omit the <code>WHENEVER SUCCESSFUL</code> clause entirely, Oracle stops auditing for all statements, regardless of success or failure.

---

## Examples

If you have chosen auditing for every SQL statement that queries the `EMP` table in the schema `SCOTT`, you can stop auditing for such queries by issuing the following statement:

```
NOAUDIT SELECT
  ON scott.emp;
```


You can stop auditing for queries that complete successfully by issuing the following statement:

```
NOAUDIT SELECT
  ON scott.emp
  WHENEVER SUCCESSFUL;
```

This statement stops auditing only for successful queries. Oracle continues to audit queries resulting in Oracle errors.

## RENAME

### Syntax



### Purpose

To rename a table, view, sequence, or private synonym for a table, view, or sequence.

- Oracle automatically transfers integrity constraints, indexes, and grants on the old object to the new object.
- Oracle invalidates all objects that depend on the renamed object, such as views, synonyms, and stored procedures and functions that refer to a renamed table.

Do not use this statement to rename public synonyms. Instead, drop the public synonym and then create another public synonym with the new name.

**See Also:** ["DROP SYNONYM"](#) on page 7-500 and ["CREATE SYNONYM"](#) on page 4-377.

### Prerequisites

The object must be in your own schema.

### Keywords and Parameters

---

<i>old</i>	is the name of an existing table, view, sequence, or private synonym.
<i>new</i>	is the new name to be given to the existing object. The new name must not already be used by another schema object in the same namespace and must follow the rules for naming schema objects defined in the section <a href="#">"Schema Object Naming Rules"</a> on page 2-72.

---

### Example

To change the name of table DEPT to EMP\_DEPT, issue the following statement:

```

    RENAME dept TO emp_dept;
  
```

You cannot use this statement directly to rename columns. However, you can rename a column using this statement together with the `CREATE TABLE` statement with `AS subquery`. The following statements re-create the table `STATIC`, renaming a column from `OLDNAME` to `NEWNAME`:


```
CREATE TABLE temporary (newname, col2, col3)
    AS SELECT oldname, col2, col3 FROM static;
```

```
DROP TABLE static;
```

```
RENAME temporary TO static;
```

## REVOKE *system\_privileges\_and\_roles*

### Syntax



### Purpose

To revoke system privileges and roles from users and roles. To revoke object privileges from users and roles, refer to ["REVOKE schema\\_object\\_privileges"](#) on page 7-559. For information on granting system privileges and roles, see ["GRANT system\\_privileges\\_and\\_roles"](#) on page 7-519.

### Prerequisites

You must have been granted the system privilege or role with the `ADMIN OPTION`. Also, you can revoke any role if you have the `GRANT ANY ROLE` system privilege.

The `REVOKE` statement can revoke only privileges and roles that were previously granted directly with a `GRANT` statement. You cannot use this statement to revoke:

- Privileges or roles not granted to the revokee
- Roles granted through the operating system
- Privileges or roles granted to the revokee through roles

### Keywords and Parameters

- 
- |                    |  |
|--------------------|--|
| <i>system_priv</i> | is a system privilege to be revoked. For a list of the system privileges, see <a href="#">Table 7-5</a> on page 7-521.   |
| ■                  | If you revoke a <b>privilege from a user</b> , Oracle removes the privilege from the user's privilege domain. Effective immediately, the user cannot exercise the privilege. |



- If you revoke **a privilege from a role**, Oracle removes the privilege from the role's privilege domain. Effective immediately, users with the role enabled cannot exercise the privilege. Also, other users who have been granted the role and subsequently enable the role cannot exercise the privilege.
- If you revoke **a privilege from PUBLIC**, Oracle removes the privilege from the privilege domain of each user who has been granted the privilege through PUBLIC. Effective immediately, such users can no longer exercise the privilege. However, the privilege is not revoked from users who have been granted the privilege directly or through roles.

**Restriction:** A system privilege cannot appear more than once in the list of privileges to be revoked.

*role*

is a role to be revoked. For a list of the roles predefined by Oracle, see "[GRANT system\\_privileges\\_and\\_roles](#)" on page 7-519.

- If you revoke **a role from a user**, Oracle makes the role unavailable to the user. If the role is currently enabled for the user, the user can continue to exercise the privileges in the role's privilege domain as long as it remains enabled. However, the user cannot subsequently enable the role.
- If you revoke **a role from another role**, Oracle removes the revoked role's privilege domain from the revokee role's privilege domain. Users who have been granted and have enabled the revokee role can continue to exercise the privileges in the revoked role's privilege domain as long as the revokee role remains enabled. However, other users who have been granted the revokee role and subsequently enable it cannot exercise the privileges in the privilege domain of the revoked role.
- If you revoke **a role from PUBLIC**, Oracle makes the role unavailable to all users who have been granted the role through PUBLIC. Any user who has enabled the role can continue to exercise the privileges in its privilege domain as long as it remains enabled. However, users cannot subsequently enable the role. The role is not revoked from users who have been granted the role directly or through other roles.

**Restriction:** A system role cannot appear more than once in the list of roles to be revoked.

FROM

identifies users and roles from which the system privileges or roles are to be revoked.

**Restriction:** A user, a role, or PUBLIC cannot appear more than once in the FROM clause.

PUBLIC

revokes the system privilege or role from all users.

## Examples

The following statement revokes the DROP ANY TABLE system privilege from the users BILL and MARY:

```
REVOKE DROP ANY TABLE
FROM bill, mary;
```

BILL and MARY can no longer drop tables in schemas other than their own.

The following statement revokes the role CONTROLLER from the user HANSON:

```
REVOKE controller
  FROM hanson;
```

HANSON can no longer enable the CONTROLLER role.

The following statement revokes the CREATE TABLESPACE system privilege from the CONTROLLER role:

```
REVOKE CREATE TABLESPACE
  FROM controller;
```

Enabling the CONTROLLER role no longer allows users to create tablespaces.

To revoke the role VP from the role CEO, issue the following statement:

```
REVOKE vp
  FROM ceo;
```

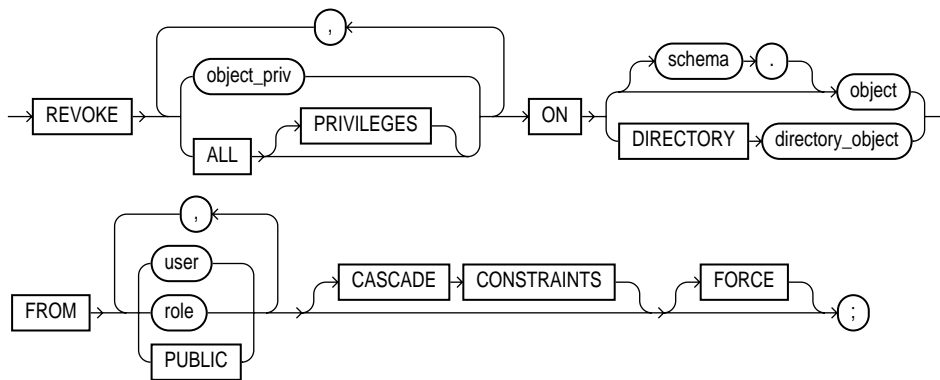
VP is no longer granted to CEO.

To revoke the CREATE ANY DIRECTORY system privilege from user SCOTT, issue the following statement:

```
REVOKE CREATE ANY DIRECTORY FROM scott;
```

## REVOKE *schema\_object\_privileges*

### Syntax



### Purpose

To revoke object privileges for a particular object from users and roles.

For information on granting schema object privileges, see ["GRANT object\\_privileges"](#) on page 7-532. To revoke system privileges or roles, refer to ["REVOKE system\\_privileges\\_and\\_roles"](#) on page 7-556.

Each object privilege authorizes some operation on an object. By revoking an object privilege, you prevent the revokee from performing that operation. However, multiple users may grant the same object privilege to the same user, role, or **PUBLIC**. To remove the privilege from the grantee's privilege domain, all grantors must revoke the privilege. If even one grantor does not revoke the privilege, the grantee can still exercise the privilege by virtue of that grant.

For a summary of the object privileges for each type of object, see [Table 7-7](#) on page 7-535.

### Prerequisites

You must have previously granted the object privileges to each user and role.

You can use the **REVOKE** statement only to revoke object privileges that you previously granted directly to the revokee. You cannot use this statement to revoke:

- Object privileges that you did not grant to the revokee
- Privileges granted through the operating system
- Privileges granted to roles granted to the revokee

## Keywords and Parameters

<i>object_priv</i>	<p>is an object privilege to be revoked. You can substitute any of the following values: ALTER, DELETE, EXECUTE, INDEX, INSERT, READ, REFERENCES, SELECT, UPDATE.</p> <ul style="list-style-type: none"> <li>■ If you revoke a <b>privilege from a user</b>, Oracle removes the privilege from the user's privilege domain. Effective immediately, the user cannot exercise the privilege. <ul style="list-style-type: none"> <li>- If that user has granted that privilege to other users or roles, Oracle also revokes the privilege from those other users or roles.</li> <li>- If that user's schema contains a procedure, function, or package that contains SQL statements that exercise the privilege, the procedure, function, or package can no longer be executed.</li> <li>- If that user's schema contains a view on that object, Oracle invalidates the view.</li> <li>- If you revoke the REFERENCES privilege from a user who has exercised the privilege to define referential integrity constraints, you must specify the CASCADE CONSTRAINTS clause.</li> </ul> </li> <li>■ If you revoke a <b>privilege from a role</b>, Oracle removes the privilege from the role's privilege domain. Effective immediately, users with the role enabled cannot exercise the privilege. Other users who have been granted the role cannot exercise the privilege after enabling the role.</li> <li>■ If you revoke a <b>privilege from PUBLIC</b>, Oracle removes the privilege from the privilege domain of each user who has been granted the privilege through PUBLIC. Effective immediately, all such users are restricted from exercising the privilege. However, the privilege is not revoked from users who have been granted the privilege directly or through roles.</li> </ul> <p><b>Restriction:</b> A privilege cannot appear more than once in the list of privileges to be revoked. A user, a role, or PUBLIC cannot appear more than once in the FROM clause.</p>
ALL PRIVILEGES	<p>revokes all object privileges that you have granted to the revokee.</p> <hr/> <p><b>Note:</b> If <i>no</i> privileges have been granted on the object, Oracle takes no action and does not return an error.</p>
ON DIRECTORY <i>directory_object</i>	<p>identifies a directory object on which privileges are revoked. You cannot qualify <i>directory_object</i> with <i>schema</i> when using the ON DIRECTORY clause. The object must be a directory. See "<a href="#">CREATE DIRECTORY</a>" on page 7-282.</p>

---

<code>ON <i>object</i></code>	<p>identifies the object on which the object privileges are revoked. This object can be</p> <ul style="list-style-type: none"> <li>■ A table, view, sequence, procedure, stored function, or package, materialized view/snapshot</li> <li>■ A synonym for a table, view, sequence, procedure, stored function, package, or materialized view/snapshot</li> <li>■ A library, indextype, or user-defined operator</li> </ul> <p>If you do not qualify <i>object</i> with <i>schema</i>, Oracle assumes the object is in your own schema.</p> <ul style="list-style-type: none"> <li>■ If you revoke the <code>SELECT</code> object privilege (with or without the <code>GRANT OPTION</code>) on the containing table or snapshot of a materialized view, the materialized view will be invalidated.</li> <li>■ If you revoke the <code>SELECT</code> object privilege (with or without the <code>GRANT OPTION</code>) on any of the master tables of a materialized view, both the view and its containing table or materialized view will be invalidated.</li> </ul>
<code>FROM</code>	<p>identifies users and roles from which the object privileges are revoked.</p> <p><code>PUBLIC</code>    revokes object privileges from all users.</p>
<code>CASCADE CONSTRAINTS</code>	<p>This clause is relevant only if you revoke the <code>REFERENCES</code> privilege or <code>ALL [PRIVILEGES]</code>. It drops any referential integrity constraints that the revokee has defined using the <code>REFERENCES</code> privilege (which might have been granted either explicitly or implicitly through a grant of <code>ALL [PRIVILEGES]</code>).</p>
<code>FORCE</code>	<p>revokes <code>EXECUTE</code> object privilege on user-defined type objects with table or type dependencies. You must use the <code>FORCE</code> clause to revoke the <code>EXECUTE</code> object privilege on user-defined type objects with table dependencies.</p> <p>If you specify <code>FORCE</code>, all privileges will be revoked, but all dependent objects are marked <code>INVALID</code>, data in dependent tables becomes inaccessible, and all dependent function-based indexes are marked <code>UNUSABLE</code>. (Regranting the necessary type privilege will revalidate the table.) For detailed information about type dependencies and user-defined object privileges, see <i>Oracle8i Concepts</i>.</p>

---

## Examples

**Basic Example** You can grant `DELETE`, `INSERT`, `SELECT`, and `UPDATE` privileges on the table `BONUS` to the user `PEDRO` with the following statement:

```
GRANT ALL
  ON bonus TO pedro;
```

To revoke the `DELETE` privilege on `BONUS` from `PEDRO`, issue the following statement:

```
REVOKE DELETE
```

```
ON bonus FROM pedro;
```

To revoke the remaining privileges on `BONUS` that you granted to `PEDRO`, issue the following statement:

```
REVOKE ALL
ON bonus FROM pedro;
```

**PUBLIC Example** You can grant `SELECT` and `UPDATE` privileges on the view `REPORTS` to all users by granting the privileges to the role `PUBLIC`:

```
GRANT SELECT, UPDATE
ON reports TO public;
```

The following statement revokes `UPDATE` privilege on `REPORTS` from all users:

```
REVOKE UPDATE
ON reports FROM public;
```

Users can no longer update the `REPORTS` view, although users can still query it. However, if you have also granted the `UPDATE` privilege on `REPORTS` to any users, either directly or through roles, these users retain the privilege.

**Schema Example** You can grant the user `BLAKE` the `SELECT` privilege on the `ESEQ` sequence in the schema `ELLY` with the following statement:

```
GRANT SELECT
ON elly.eseq TO blake;
```

To revoke the `SELECT` privilege on `ESEQ` from `BLAKE`, issue the following statement:

```
REVOKE SELECT
ON elly.eseq FROM blake;
```

However, if the user `ELLY` has also granted `SELECT` privilege on `ESEQ` to `BLAKE`, `BLAKE` can still use `ESEQ` by virtue of `ELLY`'s grant.

**CASCADE CONSTRAINTS Example** You can grant `BLAKE` the privileges `REFERENCES` and `UPDATE` on the `EMP` table in the schema `SCOTT` with the following statement:

```
GRANT REFERENCES, UPDATE
ON scott.emp TO blake;
```

BLAKE can exercise the REFERENCES privilege to define a constraint in his own DEPENDENT table that refers to the EMP table in the schema SCOTT:

```
CREATE TABLE dependent
(dependno    NUMBER,
 dependname VARCHAR2(10),
 employee   NUMBER
           CONSTRAINT in_emp REFERENCES scott.emp(ename) );
```

You can revoke the REFERENCES privilege on SCOTT.EMP from BLAKE, by issuing the following statement that contains the CASCADE CONSTRAINTS clause:

```
REVOKE REFERENCES
      ON scott.emp
      FROM blake
      CASCADE CONSTRAINTS;
```

Revoking BLAKE's REFERENCES privilege on SCOTT.EMP causes Oracle to drop the IN\_EMP constraint, because BLAKE required the privilege to define the constraint.


However, if BLAKE has also been granted the REFERENCES privilege on SCOTT.EMP by a user other than you, Oracle does not drop the constraint. BLAKE still has the privilege necessary for the constraint by virtue of the other user's grant.

**Directory Example** You can revoke READ privilege on directory BFILE\_DIR1 from SUE, by issuing the following statement:

```
REVOKE READ ON DIRECTORY bfile_dir1 FROM sue;
```

## ROLLBACK

### Syntax



### Purpose

To undo work done in the current transaction, or to manually undo the work done by an in-doubt distributed transaction.

---

---

**Note:** Oracle recommends that you explicitly end transactions in application programs using either a `COMMIT` or `ROLLBACK` statement. If you do not explicitly commit the transaction and the program terminates abnormally, Oracle rolls back the last uncommitted transaction. See also "[COMMIT](#)" on page 7-230.

---

---

#### See Also:

- *Oracle8i Concepts* for information on transactions
- "[SET TRANSACTION](#)" on page 7-602 for information on setting characteristics of the current transaction.
- "[SAVEPOINT](#)" on page 7-567.

### Prerequisites

To roll back your current transaction, no privileges are necessary.

To manually roll back an in-doubt distributed transaction that you originally committed, you must have the `FORCE TRANSACTION` system privilege. To manually roll back an in-doubt distributed transaction originally committed by another user, you must have the `FORCE ANY TRANSACTION` system privilege.



## Keywords and Parameters

---

WORK	is optional and is provided for ANSI compatibility.
TO SAVEPOINT <i>savepoint</i>	<p>rolls back the current transaction to the specified savepoint. If you omit this clause, the ROLLBACK statement rolls back the entire transaction. See also "SAVEPOINT" on page 7-567.</p> <p>Using ROLLBACK <b>without</b> the TO SAVEPOINT clause performs the following operations:</p> <ul style="list-style-type: none"> <li>■ Ends the transaction</li> <li>■ Undoes all changes in the current transaction</li> <li>■ Erases all savepoints in the transaction</li> <li>■ Releases the transaction's locks</li> </ul> <p>Using ROLLBACK <b>with</b> the TO SAVEPOINT clause performs the following operations:</p> <ul style="list-style-type: none"> <li>■ Rolls back just the portion of the transaction after the savepoint.</li> <li>■ Erases all savepoints created after that savepoint. The named savepoint is retained, so you can roll back to the same savepoint multiple times. Prior savepoints are also retained.</li> <li>■ Releases all table and row locks acquired since the savepoint. Other transactions that have requested access to rows locked after the savepoint must continue to wait until the transaction is committed or rolled back. Other transactions that have not already requested the rows can request and access the rows immediately.</li> </ul> <p><b>Restriction:</b> You cannot manually roll back an in-doubt transaction to a savepoint.</p>
FORCE	<p>manually rolls back an in-doubt distributed transaction. The transaction is identified by the 'text' containing its local or global transaction ID. To find the IDs of such transactions, query the data dictionary view DBA_2PC_PENDING.</p> <p><b>See Also:</b> <i>Oracle8i Distributed Database Systems</i> for more information on distributed transactions and rolling back in-doubt transactions.</p> <p>A ROLLBACK statement with a FORCE clause rolls back only the specified transaction. Such a statement does not affect your current transaction.</p> <p><b>Restriction:</b> ROLLBACK statements with the FORCE clause are not supported in PL/SQL.</p>

---

## Examples

The following statement rolls back your entire current transaction:

```
ROLLBACK;
```

The following statement rolls back your current transaction to savepoint SP5 :

```
ROLLBACK TO SAVEPOINT sp5;
```

The following statement manually rolls back an in-doubt distributed transaction:

```
ROLLBACK WORK  
    FORCE '25.32.87';
```

---

## SAVEPOINT

### Syntax

```
SAVEPOINT savepoint ;
```

### Purpose

To identify a point in a transaction to which you can later roll back.

#### See Also:

- *Oracle8i Concepts* for information on savepoints.
- ["ROLLBACK"](#) on page 7-564 for information on rolling back transactions.
- ["SET TRANSACTION"](#) on page 7-602 for information on setting characteristics of the current transaction

### Prerequisites

None.

### Keywords and Parameters

---

<i>savepoint</i>	is the name of the savepoint to be created.  Savepoint names must be distinct within a given transaction. If you create a second savepoint with the same identifier as an earlier savepoint, the earlier savepoint is erased. After a savepoint has been created, you can either continue processing, commit your work, roll back the entire transaction, or roll back to the savepoint.
------------------	--

---

### Example

To update BLAKE's and CLARK's salary, check that the total company salary does not exceed 27,000, then reenter CLARK's salary, enter:

```
UPDATE emp
  SET sal = 2000
  WHERE ename = 'BLAKE';
SAVEPOINT blake_sal;
```

## SAVEPOINT

---

```
UPDATE emp
  SET sal = 1500
  WHERE ename = 'CLARK';
SAVEPOINT clark_sal;

SELECT SUM(sal) FROM emp;

ROLLBACK TO SAVEPOINT blake_sal;

UPDATE emp
  SET sal = 1200
  WHERE ename = 'CLARK';

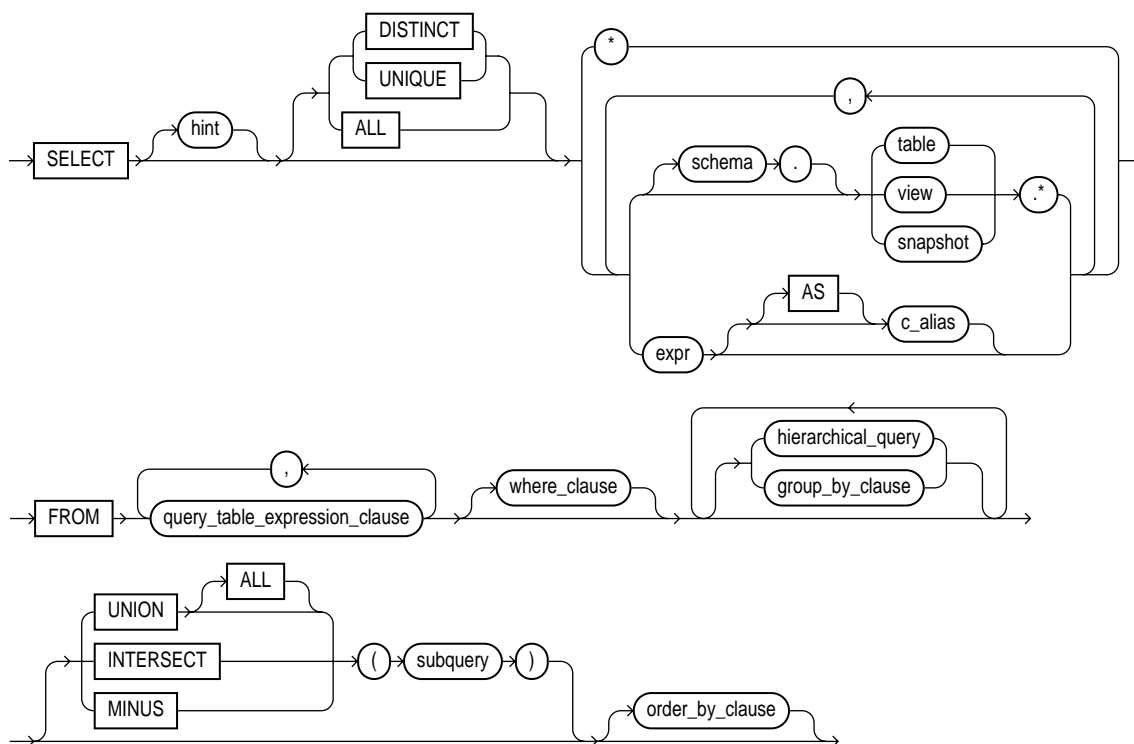
COMMIT;
```

## SELECT and Subqueries

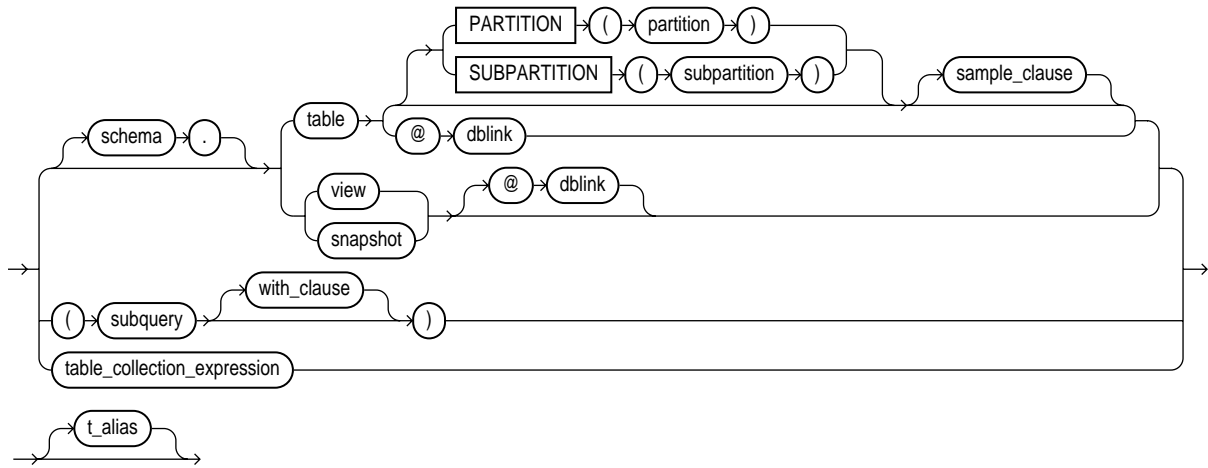
### Syntax




**subquery::=**




**query\_table\_expression\_clause::=**




**sample\_clause::=**









**with\_clause::=**



**table\_collection\_expression::=**



**where\_clause::=****outer\_join::=****hierarchical\_query\_clause::=****group\_by\_clause::=****order\_by\_clause::=**

**for\_update\_clause::=****Purpose**

To retrieve data from one or more tables, object tables, views, object views, or materialized views. For general information on queries and subqueries, see ["Queries and Subqueries"](#) on page 5-19.

---



---

**Note:** If the result (or part of the result) of a `SELECT` statement is equivalent to an existing materialized view, Oracle may use the materialized view in place of one or more tables specified in the `SELECT` statement. This substitution is called **query rewrite**, and takes place only if cost optimization is enabled and the `QUERY_REWRITE_ENABLED` parameter is set to `TRUE`. To determine whether query write has occurred, use the `EXPLAIN PLAN` statement (see ["EXPLAIN PLAN"](#) on page 7-512). For more information on materialized views and query rewrite, see *Oracle8i Data Warehousing Guide*.

---



---

**Prerequisites**

For you to select data from a table or materialized view, the table or materialized view must be in your own schema or you must have the `SELECT` privilege on the table or materialized view.

For you to select rows from the base tables of a view,

- You must have the `SELECT` privilege on the view, and
- Whoever owns the schema containing the view must have the `SELECT` privilege on the base tables.

The `SELECT ANY TABLE` system privilege also allows you to select data from any table or any materialized view or any view's base table.



## Keywords and Parameters

---

<i>hint</i>	is a comment that passes instructions to the optimizer on choosing an execution plan for the statement. For the syntax and description of hints, see <a href="#">"Hints"</a> on page 2-63 and <i>Oracle8i Designing and Tuning for Performance</i> .
DISTINCT   UNIQUE	returns only one copy of each set of duplicate rows selected (these two keywords are synonymous). Duplicate rows are those with matching values for each expression in the select list.  <b>Restrictions:</b> <ul style="list-style-type: none"> <li>■ When you specify DISTINCT or UNIQUE, the total number of bytes in all select list expressions is limited to the size of a data block minus some overhead. This size is specified by the initialization parameter DB_BLOCK_SIZE.</li> <li>■ You cannot specify DISTINCT if the FROM clause contains LOB columns.</li> </ul>
ALL	returns all rows selected, including all copies of duplicates. The default is ALL.
*	selects all columns from all tables, views, or materialized views listed in the FROM clause.  <b>Note:</b> If you are selecting from a table (that is, you specify <i>table</i> in the FROM clause rather than an <i>view</i> or a <i>materialized view</i> ), columns that have been marked as UNUSED by the ALTER TABLE SET UNUSED statement are not selected.  <b>See Also:</b> <a href="#">"ALTER TABLE"</a> on page 7-123.
<i>schema</i>	is the schema containing the selected table, view, or materialized view. If you omit <i>schema</i> , Oracle assumes the table, view, or materialized view is in your own schema.
<i>table.*</i>   <i>view.*</i>   <i>snapshot.*</i>	selects all columns from the specified table, view, or materialized view. You can use the schema qualifier to select from a table, view, or materialized view in a schema other than your own. A query that selects rows from two or more tables, views, or materialized views is a join.  <b>See Also:</b> <a href="#">"Joins"</a> on page 5-22.
<i>expr</i>	selects an expression. A column name in this list can be qualified with <i>schema</i> only if the table, view, or materialized view containing the column is qualified with <i>schema</i> in the FROM clause.  <b>See Also:</b> the syntax description of <i>expr</i> in <a href="#">"Expressions"</a> on page 5-1.

**Restrictions:**

- If you also specify a *group\_by\_clause* in this statement, this select list can contain only the following types of expressions:
  - constants
  - aggregate functions and the functions `USER`, `UID`, and `SYSDATE`
  - expressions identical to those in the *group\_by\_clause*
  - expressions involving the above expressions that evaluate to the same value for all rows in a group
- You can select a rowid from a join view only if the join has one and only one key-preserved table. The rowid of that table becomes the rowid of the view. For information on key-preserved tables, see *Oracle8i Administrator's Guide*.
- If two or more tables have some column names in common, you must qualify column names with names of tables.

*c\_alias* provides a different name for the column expression and causes the alias to be used in the column heading. The `AS` keyword is optional. The alias effectively renames the select list item for the duration of the query. The alias can be used in the *order\_by\_clause*, but not other clauses in the query.

`FROM` specifies the table, view, materialized view, or partition from which data is selected, or a subquery that specifies the objects from which data is selected.

*query\_table\_expression\_clause*

- sample\_clause* causes Oracle to select from a random sample of rows from the table, rather than from the entire table.
- `BLOCK` instructs Oracle to perform random block sampling instead of random row sampling. For a discussion of the difference, refer to *Oracle8i Concepts*.
  - *sample\_percent* is a number specifying the percentage of the total row or block count to be included in the sample. The value must be in the range .000001 to (but not including) 100.

**Restrictions:**

- You can specify `SAMPLE` only in a query that selects from a single table. Joins are not supported. However, you can achieve the same results by using a `CREATE TABLE ... AS SELECT` query to materialize a sample of an underlying table and then rewrite the original query to refer to the newly created table sample. If you wish, you can write additional queries to materialize samples for other tables.

**See Also:** "[SAMPLE Examples](#):" on page 7-582.

- When you specify `SAMPLE`, Oracle automatically uses cost-based optimization. Rule-based optimization is not supported with this clause.

**WARNING:** The use of statistically incorrect assumptions when using this feature can lead to incorrect or undesirable results. Refer to *Oracle8i Concepts* for more information on using the `sample_clause`.

`PARTITION`  
(*partition*)

specifies partition-level data retrieval. The *partition* parameter may be the name of the partition within *table* from which to retrieve data or a more complicated predicate restricting retrieval to just one partition of the table.

`SUBPARTITION`  
(*subpartition*)

*dblink*

is the complete or partial name for a database link to a remote database where the table, view, or materialized view is located. This database need not be an Oracle database.

**See Also:**

- "[Referring to Objects in Remote Databases](#)" on page 2-79 for more information on referring to database links.
- "[Distributed Queries](#)" on page 5-26 for more information about distributed queries.

If you omit *dblink*, Oracle assumes that the table, view, or materialized view is on the local database.

**Restriction:** You cannot query a user-defined type or an object REF on a remote table.

*table, view,*  
*snapshot*

is the name of a table, view, or materialized view from which data is selected. "Materialized view" is synonymous with "snapshot".

*with\_clause*

restricts the subquery in one of the following ways:

- `WITH READ ONLY` specifies that the subquery cannot be updated.

- `WITH CHECK OPTION` specifies that, if the subquery is used in place of a table in an `INSERT`, `UPDATE`, or `DELETE` statement, Oracle prohibits any changes to that table that would produce rows that are not included in the subquery.

**See Also:** the [WITH CHECK OPTION Example](#) on page 7-587.

*table\_collection\_expression* informs Oracle that the collection value expression should be treated as a table for purposes of query and DML operations. The *collection\_expression* can be a subquery, a column, a `CAST` or `DECODE` expression, a function, or a collection constructor. Regardless of its form, it must return a collection value (that is, a value whose type is nested table or varray). This process of extracting the elements of a collection is called **collection unnesting**.

**See Also:** "[Collection Unnesting Examples](#)" on page 7-593.

The *collection\_expression* can reference columns of tables defined to its left in the `FROM` clause. This is called **left correlation**. Left correlation can occur only in *table\_collection\_expression*. Other subqueries cannot contain references to columns defined outside the subquery.

The optional `(+)` lets you specify that *table\_collection\_expression* should return a row with all fields set to `NULL` if the collection is null or empty. The `(+)` is valid only if *collection\_expression* uses left correlation. The result is similar to that of an outer join.

**See Also:** "[Outer Joins](#)" on page 5-23.

---

**Note:** In earlier releases of Oracle, when *collection\_expression* was a subquery, *table\_collection\_expr* was expressed as "THE subquery". That usage is now deprecated.

---

*t\_alias* provides a **correlation name** for the table, view, materialized view, or subquery for evaluating the query and is most often used in a correlated query. Other references to the table, view, or materialized view throughout the query must refer to the alias.

---

**Note:** This alias is **required** if the *query\_table\_expression\_clause* references any object type attributes or object type methods.

---

---

<i>where_clause</i>	<p>restricts the rows selected to those that satisfy one or more conditions.</p> <ul style="list-style-type: none"> <li>■ <i>condition</i> can be any valid SQL condition. <b>See Also:</b> the syntax description of condition in "Conditions" on page 5-14.</li> <li>■ <i>outer_join</i> applies only if the <i>query_table_expression_clause</i> specifies more than one table. This special form of condition requires Oracle to return all the rows that satisfy the condition, as well as all the rows from one of the tables for which no rows of the other table satisfy the condition.  If one of the elements in the <i>query_table_expression_clause</i> is actually a nested table or some other form of collection, you specify the outer-join syntax in the <i>table_collection_expression</i> rather than in the <i>where_clause</i>.  <b>See Also:</b> "Outer Joins" on page 5-23 for more information, including rules and restrictions that apply to outer joins.</li> </ul> <p>If you omit this clause, Oracle returns all rows from the tables, views, or materialized views in the FROM clause.</p> <hr/> <p><b>Note:</b> If this clause refers to a DATE column of a partitioned table or index, Oracle performs partition pruning only if (1) you created the table or index partitions by fully specifying the year using the TO_DATE function with a 4-digit format mask, <b>and</b> (2) you specify the date in the query's <i>where_clause</i> using the TO_DATE function and either a 2- or 4-digit format mask.</p> <hr/> <p><b>See Also:</b> the "PARTITION Example" on page 7-582.</p>
<i>hierarchical_query_clause</i>	<p>lets you select rows in a hierarchical order. For a discussion of hierarchical queries, see "Hierarchical Queries" on page 5-20.</p> <p>The preceding <i>where_clause</i>, if specified, restricts the rows returned by the query without affecting other rows of the hierarchy.</p> <p>SELECT statements that contain hierarchical queries can contain the LEVEL pseudocolumn. LEVEL returns the value 1 for a root node, 2 for a child node of a root node, 3 for a grandchild, etc. The number of levels returned by a hierarchical query may be limited by available user memory.</p> <p><b>See Also:</b> "Pseudocolumns" on page 2-56 for more information on LEVEL.</p> <p><b>Restrictions:</b> If you specify a hierarchical query:</p> <ul style="list-style-type: none"> <li>■ The same statement cannot also perform a join.</li> <li>■ The same statement cannot also select data from a view whose query performs a join.</li> <li>■ If you also specify the <i>order_by_clause</i>, it takes precedence over any ordering specified by the hierarchical query.</li> </ul> <p>START WITH identifies the row(s) to be used as the root(s) of a hierarchical query. This clause specifies a condition that the roots must satisfy. If you omit this clause, Oracle uses all rows in the table as root rows. The START WITH <i>condition</i> can contain a subquery.</p>

---

CONNECT BY	<p>specifies the relationship between parent rows and child rows of the hierarchy. <i>condition</i> can be any condition as described in "<a href="#">Conditions</a>" on page 5-14. However, some part of the condition must use the PRIOR operator to refer to the parent row. The part of the condition containing the PRIOR operator must have one of the following forms:</p> <pre style="margin-left: 40px;">PRIOR expr comparison_operator expr expr comparison_operator PRIOR expr</pre> <p><b>Restriction:</b> The CONNECT BY condition cannot contain a subquery.</p>
<i>group_by_clause</i>	<p>groups the selected rows based on the value of <i>expr(s)</i> for each row, and returns a single row of summary information for each group. If this clause contains CUBE or ROLLUP extensions, then superaggregate groupings are produced in addition to the regular groupings.</p> <p>Expressions in the <i>group_by_clause</i> can contain any columns in the tables, views, and materialized views in the FROM clause, regardless of whether the columns appear in the select list.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"> <li>■ The <i>group_by_clause</i> can contain no more than 255 expressions.</li> <li>■ You cannot specify LOB columns, nested tables, or varrays as part of <i>expr</i>.</li> <li>■ The total number of bytes in all expressions in the <i>group_by_clause</i> is limited to the size of a data block (specified by the initialization parameter DB_BLOCK_SIZE) minus some overhead.</li> <li>■ If the <i>group_by_clause</i> references any object columns, the query will not be parallelized.</li> </ul>
ROLLUP	<p>is an extension to the <i>group_by_clause</i> that groups the selected rows based on the values of the first <i>n</i>, <i>n</i>-1, <i>n</i>-2, ... 0 expressions for each row, and returns a single row of summary for each group. You can use the ROLLUP operation to produce <b>subtotal values</b>.</p> <p>For example, given three expressions in the ROLLUP clause of the <i>group_by_clause</i>, the operation results in <math>n+1 = 3+1 = 4</math> groupings.</p> <p>Rows based on the values of the first 'n' expressions are called <b>regular rows</b>, and the others are called <b>superaggregate rows</b>.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"> <li>■ "<a href="#">GROUPING</a>" on page 4-36 for an example</li> <li>■ <i>Oracle8i Data Warehousing Guide</i>.</li> </ul>
CUBE	<p>is an extension to the <i>group_by_clause</i> that groups the selected rows based on the values of all possible combinations of expressions for each row, and returns a single row of summary information for each group. You can use the CUBE operation to produce <b>cross-tabulation values</b>.</p>

For example, given three expressions in the `CUBE` clause of the `group_by_clause`, the operation results in  $2^n = 2^3 = 8$  groupings. Rows based on the values of 'n' expressions are called regular rows, and the rest are called **superaggregate rows**.

**See Also:**

- ["GROUPING"](#) on page 4-36
- ["CUBE Example"](#) on page 7-583 for an example
- *Oracle8i Data Warehousing Guide*.

HAVING

restricts the groups of rows returned to those groups for which the specified *condition* is `TRUE`. If you omit this clause, Oracle returns summary rows for all groups.

Specify `GROUP BY` and `HAVING` after the *where\_clause* and `CONNECT BY` clause. If you specify both `GROUP BY` and `HAVING`, they can appear in either order.

**See Also:** the syntax description of *expr* in ["Expressions"](#) on page 5-1 and the syntax description of *condition* in ["Conditions"](#) on page 5-14.

UNION | UNION  
ALL |  
INTERSECT |  
MINUS

are set operators that combine the rows returned by two `SELECT` statements into a single result. The number and datatypes of the columns selected by each component query must be the same, but the column lengths can be different.

If you combine more than two queries with set operators, Oracle evaluates adjacent queries from left to right. You can use parentheses to specify a different order of evaluation.

For information on these operators, see ["Set Operators"](#) on page 3-12.

**Restrictions:**

- These set operators are not valid on columns of type `BLOB`, `CLOB`, `BFILE`, `varray`, or nested table.
- The `UNION`, `INTERSECT`, and `MINUS` operators are not valid on `LONG` columns.
- To reference a column, you must use an alias to name the column.
- You cannot also specify the *for\_update\_clause* with these set operators.
- You cannot specify the *order\_by\_clause* in the *subquery* of these operators.
- You cannot use these operators in `SELECT` statements containing `TABLE` collection expressions.
- The total number of bytes in all select list expressions of a component query is limited to the size of a data block (specified by the initialization parameter `DB_BLOCK_SIZE`) minus some overhead.

**Note:** To comply with emerging SQL standards, a future release of Oracle will give the `INTERSECT` operator greater precedence than the other set operators. Therefore, you should use parentheses to specify order of evaluation in queries that use the `INTERSECT` operator with other set operators.

*order\_by\_clause* orders rows returned by the statement. Without an *order\_by\_clause*, no guarantee exists that the same query executed more than once will retrieve rows in the same order. For a discussion of ordering query results, see "[Sorting Query Results](#)" on page 5-21.

- *expr* orders rows based on their value for *expr*. The expression is based on columns in the select list or columns in the tables, views, or materialized views in the FROM clause.
- *position* orders rows based on their value for the expression in this position of the select list; *position* must be an integer.

You can specify multiple expressions in the *order\_by\_clause*. Oracle first sorts rows based on their values for the first expression. Rows with the same value for the first expression are then sorted based on their values for the second expression, and so on. Oracle sorts nulls following all others in ascending order and preceding all others in descending order.

ASC | DESC specifies the ordering sequence. ASC is the default.

NULLS FIRST | NULLS LAST specifies whether returned rows containing null values should appear first or last in the ordering sequence.

NULLS LAST is the default for ascending order, and NULLS FIRST is the default for descending order.

**Restrictions:**

- If you have specified the DISTINCT operator in this statement, this clause cannot refer to columns unless they appear in the select list.
- An *order\_by\_clause* can contain no more than 255 expressions.
- You cannot order by a LOB column, nested table, or varray.

If you specify a *group\_by\_clause* in the same statement, this *order\_by\_clause* is restricted to the following expressions:

- Constants
- Aggregate functions
- Analytic functions
- The functions USER, UID, and SYSDATE
- Expressions identical to those in the *group\_by\_clause*
- Expressions involving the above expressions that evaluate to the same value for all rows in a group.

*for\_update\_clause* locks the selected rows so that other users cannot lock or update the rows until you end your transaction. You can specify this clause only in a top-level SELECT statement (not in subqueries).



- Prior to updating a LOB value, you must lock the row containing the LOB. One way to lock the row is with a `SELECT ... FOR UPDATE` statement.

**See Also:** "LOB Locking Example" on page 7-585.

- Nested table rows are not locked as a result of locking the parent table rows. If you want the nested table rows to be locked, you must lock them explicitly.

**OF** Locks the select rows only for a particular table in a join. The columns in the OF clause only specify which tables' rows are locked. The specific columns of the table that you specify are not significant. If you omit this clause, Oracle locks the selected rows from all the tables in the query.

**NOWAIT** returns control to you if the `SELECT` statement attempts to lock a row that is locked by another user. If you omit this clause, Oracle waits until the row is available and then returns the results of the `SELECT` statement.

#### Restrictions:

- You cannot specify this clause with the following other constructs: the `DISTINCT` or `CURSOR` operator, set operators, *group\_by\_clause*, or aggregate functions.
- The tables locked by this clause must all be located on the same database, and on the same database as any `LONG` columns and sequences referenced in the same statement.

## Examples

**Simple Query Examples** The following statement selects rows from the `EMP` table with the department number of 30:

```
SELECT *
   FROM emp
  WHERE deptno = 30;
```

The following statement selects the name, job, salary and department number of all employees except sales people from department number 30:

```
SELECT ename, job, sal, deptno
   FROM emp
  WHERE NOT (job = 'SALESMAN' AND deptno = 30);
```

The following statement selects from subqueries in the `FROM` clause and gives departments' total employees and salaries as a decimal value of all the departments:

```
SELECT a.deptno "Department",
       a.num_emp/b.total_count "%Employees",
       a.sal_sum/b.total_sal  "%Salary"
   FROM
```

```
(SELECT deptno, COUNT(*) num_emp, SUM(SAL) sal_sum
  FROM scott.emp
  GROUP BY deptno) a,
(SELECT COUNT(*) total_count, SUM(sal) total_sal
  FROM scott.emp) b ;
```

**PARTITION Example** You can select rows from a single partition of a partitioned table by specifying the keyword `PARTITION` in the `FROM` clause. This SQL statement assigns an alias for and retrieves rows from the `NOV98` partition of the `SALES` table:

```
SELECT * FROM sales PARTITION (nov98) s
  WHERE s.amount_of_sale > 1000;
```

The following example selects rows from the `SALES` table for sales earlier than a specified date:

```
SELECT * FROM sales
  WHERE sale_date < TO_DATE('1998-06-15', 'YYYY-MM-DD');
```

**SAMPLE Examples:** The following query estimates the number of employees in the `EMP` table:

```
SELECT COUNT(*) * 100 FROM emp SAMPLE BLOCK (1);
```

The following example creates a sampled subset of the `EMP` table and then joins the resulting sampled table with `DEPT`. This operation circumvents the restriction that you cannot specify the *sample\_clause* in join queries:

```
CREATE TABLE sample_emp AS SELECT empno, deptno FROM emp SAMPLE(10);
SELECT e.empno FROM sample_emp e, dept d
  WHERE e.deptno = d.deptno AND d.name = 'DEV';
```

**GROUP BY Examples** To return the minimum and maximum salaries for each department in the employee table, issue the following statement:

```
SELECT deptno, MIN(sal), MAX (sal)
  FROM emp
  GROUP BY deptno;
```

DEPTNO	MIN(SAL)	MAX(SAL)
10	1300	5000
20	800	3000
30	950	2850

To return the minimum and maximum salaries for the clerks in each department, issue the following statement:

```
SELECT deptno, MIN(sal), MAX (sal)
   FROM emp
   WHERE job = 'CLERK'
   GROUP BY deptno;
```

DEPTNO	MIN(SAL)	MAX(SAL)
10	1300	1300
20	800	1100
30	950	950

**CUBE Example** To return the number of employees and their average yearly salary across all possible combinations of department and job category, issue the following query:

```
SELECT DECODE(GROUPING(dname), 1, 'All Departments',
             dname) AS dname,
       DECODE(GROUPING(job), 1, 'All Jobs', job) AS job,
       COUNT(*) "Total Empl", AVG(sal) * 12 "Average Sal"
   FROM emp, dept
   WHERE dept.deptno = emp.deptno
   GROUP BY CUBE (dname, job);
```

DNAME	JOB	Total Empl	Average Sa
ACCOUNTING	CLERK	1	15600
ACCOUNTING	MANAGER	1	29400
ACCOUNTING	PRESIDENT	1	60000
ACCOUNTING	All Jobs	3	35000
RESEARCH	ANALYST	2	36000
RESEARCH	CLERK	2	11400
RESEARCH	MANAGER	1	35700
RESEARCH	All Jobs	5	26100
SALES	CLERK	1	11400
SALES	MANAGER	1	34200
SALES	SALESMAN	4	16800
SALES	All Jobs	6	18800
All Departments	ANALYST	2	36000
All Departments	CLERK	4	12450
All Departments	MANAGER	3	33100
All Departments	PRESIDENT	1	60000
All Departments	SALESMAN	4	16800

All Departments All Jobs

14 24878.5714

**Hierarchical Query Examples** The following `CONNECT BY` clause defines a hierarchical relationship in which the `EMPNO` value of the parent row is equal to the `MGR` value of the child row:

```
CONNECT BY PRIOR empno = mgr;
```

In the following `CONNECT BY` clause, the `PRIOR` operator applies only to the `EMPNO` value. To evaluate this condition, Oracle evaluates `EMPNO` values for the parent row and `MGR`, `SAL`, and `COMM` values for the child row:

```
CONNECT BY PRIOR empno = mgr AND sal > comm;
```

To qualify as a child row, a row must have a `MGR` value equal to the `EMPNO` value of the parent row and it must have a `SAL` value greater than its `COMM` value.

**HAVING Example** To return the minimum and maximum salaries for the clerks in each department whose lowest salary is below \$1,000, issue the next statement:

```
SELECT deptno, MIN(sal), MAX (sal)
   FROM emp
  WHERE job = 'CLERK'
  GROUP BY deptno
  HAVING MIN(sal) < 1000;
```

DEPTNO	MIN(SAL)	MAX(SAL)
20	800	1100
30	950	950

**ORDER BY Examples** To select all salesmen's records from `EMP`, and order the results by commission in descending order, issue the following statement:

```
SELECT *
   FROM emp
  WHERE job = 'SALESMAN'
  ORDER BY comm DESC;
```

To select the employees from `EMP` ordered first by ascending department number and then by descending salary, issue the following statement:

```
SELECT ename, deptno, sal
   FROM emp
  ORDER BY deptno ASC, sal DESC;
```

To select the same information as the previous `SELECT` and use the positional `ORDER BY` notation, issue the following statement:

```
SELECT ename, deptno, sal
       FROM emp
       ORDER BY 2 ASC, 3 DESC;
```

**FOR UPDATE Examples** The following statement locks rows in the `EMP` table with clerks located in New York and locks rows in the `DEPT` table with departments in New York that have clerks:

```
SELECT empno, sal, comm
       FROM emp, dept
       WHERE job = 'CLERK'
             AND emp.deptno = dept.deptno
             AND loc = 'NEW YORK'
       FOR UPDATE;
```

The following statement locks only those rows in the `EMP` table with clerks located in New York. No rows are locked in the `DEPT` table:

```
SELECT empno, sal, comm
       FROM emp, dept
       WHERE job = 'CLERK'
             AND emp.deptno = dept.deptno
             AND loc = 'NEW YORK'
       FOR UPDATE OF emp.sal;
```

**LOB Locking Example** The following example uses a `SELECT ... FOR UPDATE` statement to lock a row containing a LOB prior to updating the LOB value.

```
INSERT INTO t_table VALUES (1, 'abcd');
```

```
COMMIT;
```

```
DECLARE
```

```
  num_var      NUMBER;
  clob_var     CLOB;
  clob_locked  CLOB;
  write_amount NUMBER;
  write_offset NUMBER;
  buffer       VARCHAR2(20) := 'efg';
```

```
BEGIN
```

```
  SELECT clob_col INTO clob_locked FROM t_table
 WHERE num_col = 1 FOR UPDATE;
```

```
write_amount := 3;
dbms_lob.write(clob_locked, write_amount, write_offset, buffer);
END;
```

**Table Collection Examples** You can perform DML operations on nested tables only if they are defined as columns of a table. Therefore, when the *query\_table\_expression\_clause* of an INSERT, DELETE, or UPDATE statement is a *table\_collection\_expression*, the collection expression must be a subquery that selects the table's nested table column. The examples that follow are based on this scenario:

```
CREATE TYPE ProjectType AS OBJECT(
    pno    NUMBER,
    pname  CHAR(31),
    budget NUMBER);
CREATE TYPE ProjectSet AS TABLE OF ProjectType;

CREATE TABLE Dept (dno NUMBER, dname CHAR(31), projs ProjectSet)
    NESTED TABLE projs STORE AS
        ProjectSetTable ((Primary Key(Nested_Table_Id, pno))
    ORGANIZATION
    INDEX COMPRESS 1);

INSERT INTO Dept VALUES (1, 'Engineering', ProjectSet());
```

This example inserts into the 'Engineering' department's 'projs' nested table:

```
INSERT INTO TABLE(SELECT d.projs
    FROM    Dept d
    WHERE   d.dno = 1)
VALUES (1, 'Collection Enhancements', 10000);
```

This example updates the 'Engineering' department's 'projs' nested table:

```
UPDATE TABLE(SELECT d.projs
    FROM    Dept d
    WHERE   d.dno = 1) p
SET p.budget = p.budget + 1000;
```

This example deletes from the 'Engineering' department's 'projs' nested table

```
DELETE TABLE(SELECT d.projs
    FROM    Dept d
    WHERE   d.dno = 1) p
WHERE p.budget > 100000;
```

**Subquery Examples** To determine who works in Taylor's department, issue the following statement:

```
SELECT ename, deptno
   FROM emp
  WHERE deptno =
      (SELECT deptno
       FROM emp
       WHERE ename = 'TAYLOR');
```

To give all employees in the EMP table a 10% raise if they have not already been issued a bonus (if they do not appear in the BONUS table), issue the following statement:

```
UPDATE emp
   SET sal = sal * 1.1
   WHERE empno NOT IN (SELECT empno FROM bonus);
```

To create a duplicate of the DEPT table named NEWDEPT, issue the following statement:

```
CREATE TABLE newdept (deptno, dname, loc)
   AS SELECT deptno, dname, loc FROM dept;
```

**WITH CHECK OPTION Example** The following statement is legal even though the second value violates the condition of the subquery *where\_clause*:

```
INSERT INTO
  (SELECT ename, deptno FROM emp WHERE deptno < 10)
  VALUES ('Taylor', 20);
```

However, the following statement is illegal because of the WITH CHECK OPTION clause:

```
INSERT INTO
  (SELECT ename, deptno FROM emp
   WHERE deptno < 10
   WITH CHECK OPTION)
  VALUES ('Taylor', 20);
```

**Equijoin Examples** This equijoin returns the name and job of each employee and the number and name of the department in which the employee works:

```
SELECT ename, job, dept.deptno, dname
   FROM emp, dept
  WHERE emp.deptno = dept.deptno;
```

ENAME	JOB	DEPTNO	DNAME
CLARK	MANAGER	10	ACCOUNTING
KING	PRESIDENT	10	ACCOUNTING
MILLER	CLERK	10	ACCOUNTING
SMITH	CLERK	20	RESEARCH
ADAMS	CLERK	20	RESEARCH
FORD	ANALYST	20	RESEARCH
SCOTT	ANALYST	20	RESEARCH
JONES	MANAGER	20	RESEARCH
ALLEN	SALESMAN	30	SALES
BLAKE	MANAGER	30	SALES
MARTIN	SALESMAN	30	SALES
JAMES	CLERK	30	SALES
TURNER	SALESMAN	30	SALES
WARD	SALESMAN	30	SALES

You must use a join to return this data because employee names and jobs are stored in a different table than department names. Oracle combines rows of the two tables according to this join condition:

```
emp.deptno = dept.deptno
```

The following equijoin returns the name, job, department number, and department name of all clerks:

```
SELECT ename, job, dept.deptno, dname
       FROM emp, dept
       WHERE emp.deptno = dept.deptno
       AND job = 'CLERK';
```

ENAME	JOB	DEPTNO	DNAME
MILLER	CLERK	10	ACCOUNTING
SMITH	CLERK	20	RESEARCH
ADAMS	CLERK	20	RESEARCH
JAMES	CLERK	30	SALES

This query is identical to the preceding example, except that it uses an additional *where\_clause* condition to return only rows with a `JOB` value of 'CLERK'.

**Self Join Example** The following query uses a self join to return the name of each employee along with the name of the employee's manager:



```
SELECT e1.ename||' works for '||e2.ename
"Employees and their Managers"
      FROM emp e1, emp e2   WHERE e1.mgr = e2.empno;
```

Employees and their Managers

```
-----
BLAKE works for KING
CLARK works for KING
JONES works for KING
FORD works for JONES
SMITH works for FORD
ALLEN works for BLAKE
WARD works for BLAKE
MARTIN works for BLAKE
SCOTT works for JONES
TURNER works for BLAKE
ADAMS works for SCOTT
JAMES works for BLAKE
MILLER works for CLARK
```

The join condition for this query uses the aliases E1 and E2 for the EMP table:

```
e1.mgr = e2.empno
```

**Outer Join Examples** This query uses an outer join to extend the results of the Equijoin example above:

```
SELECT ename, job, dept.deptno, dname
      FROM emp, dept
      WHERE emp.deptno (+) = dept.deptno;
```

ENAME	JOB	DEPTNO	DNAME
CLARK	MANAGER	10	ACCOUNTING
KING	PRESIDENT	10	ACCOUNTING
MILLER	CLERK	10	ACCOUNTING
SMITH	CLERK	20	RESEARCH
ADAMS	CLERK	20	RESEARCH
FORD	ANALYST	20	RESEARCH
SCOTT	ANALYST	20	RESEARCH
JONES	MANAGER	20	RESEARCH
ALLEN	SALESMAN	30	SALES
BLAKE	MANAGER	30	SALES
MARTIN	SALESMAN	30	SALES
JAMES	CLERK	30	SALES

```

TURNER      SALESMAN      30 SALES
WARD        SALESMAN      30 SALES
                                40 OPERATIONS
    
```

In this outer join, Oracle returns a row containing the OPERATIONS department even though no employees work in this department. Oracle returns NULL in the ENAME and JOB columns for this row. The join query in this example selects only departments that have employees.

The following query uses an outer join to extend the results of the preceding example:

```

SELECT ename, job, dept.deptno, dname
       FROM emp, dept
       WHERE emp.deptno (+) = dept.deptno
              AND job (+) = 'CLERK';
    
```

```

ENAME      JOB      DEPTNO      DNAME
-----
MILLER     CLERK      10 ACCOUNTING
SMITH      CLERK      20 RESEARCH
ADAMS      CLERK      20 RESEARCH
JAMES      CLERK      30 SALES
                                40 OPERATIONS
    
```

In this outer join, Oracle returns a row containing the OPERATIONS department even though no clerks work in this department. The (+) operator on the JOB column ensures that rows for which the JOB column is NULL are also returned. If this (+) were omitted, the row containing the OPERATIONS department would not be returned because its JOB value is not 'CLERK'.

This example shows four outer join queries on the CUSTOMERS, ORDERS, LINEITEMS, and PARTS tables. These tables are shown here:

```

SELECT custno, custname
       FROM customers;
    
```

```

CUSTNO      CUSTNAME
-----
           1 Angelic Co.
           2 Believable Co.
           3 Cables R Us
    
```

```

SELECT orderno, custno,
       TO_CHAR(orderdate, 'MON-DD-YYYY') "ORDERDATE"
       FROM orders;
    
```

ORDERNO	CUSTNO	ORDERDATE
9001	1	OCT-13-1998
9002	2	OCT-13-1998
9003	1	OCT-20-1998
9004	1	OCT-27-1998
9005	2	OCT-31-1998

```
SELECT orderno, lineno, partno, quantity
FROM lineitems;
```

ORDERNO	LINENO	PARTNO	QUANTITY
9001	1	101	15
9001	2	102	10
9002	1	101	25
9002	2	103	50
9003	1	101	15
9004	1	102	10
9004	2	103	20

```
SELECT partno, partname
FROM parts;
```

PARTNO	PARTNAME
101	X-Ray Screen
102	Yellow Bag
103	Zoot Suit

The customer Cables R Us has placed no orders, and order number 9005 has no line items.

The following outer join returns all customers and the dates they placed orders. The (+) operator ensures that customers who placed no orders are also returned:

```
SELECT custname, TO_CHAR(orderdate, 'MON-DD-YYYY') "ORDERDATE"
FROM customers, orders
WHERE customers.custno = orders.custno (+);
```

CUSTNAME	ORDERDATE
Angelic Co.	OCT-13-1993
Angelic Co.	OCT-20-1993

```

Angelic Co.          OCT-27-1993
Believable Co.     OCT-13-1993
Believable Co.     OCT-31-1993
Cables R Us

```

The following outer join builds on the result of the previous one by adding the `LINEITEMS` table to the `FROM` clause, columns from this table to the select list, and a join condition joining this table to the `ORDERS` table to the *where\_clause*. This query joins the results of the previous query to the `LINEITEMS` table and returns all customers, the dates they placed orders, and the part number and quantity of each part they ordered. The first (+) operator serves the same purpose as in the previous query. The second (+) operator ensures that orders with no line items are also returned:

```

SELECT custname,
       TO_CHAR(orderdate, 'MON-DD-YYYY') "ORDERDATE",
       partno,
       quantity
       FROM customers, orders, lineitems
       WHERE customers.custno = orders.custno (+)
       AND orders.orderno = lineitems.orderno (+);

```

CUSTNAME	ORDERDATE	PARTNO	QUANTITY
Angelic Co.	OCT-13-1993	101	15
Angelic Co.	OCT-13-1993	102	10
Angelic Co.	OCT-20-1993	101	15
Angelic Co.	OCT-27-1993	102	10
Angelic Co.	OCT-27-1993	103	20
Believable Co.	OCT-13-1993	101	25
Believable Co.	OCT-13-1993	103	50
Believable Co.	OCT-31-1993		
Cables R Us			

The following outer join builds on the result of the previous one by adding the `PARTS` table to the `FROM` clause, the `PARTNAME` column from this table to the select list, and a join condition joining this table to the `LINEITEMS` table to the *where\_clause*. This query joins the results of the previous query to the `PARTS` table to return all customers, the dates they placed orders, and the quantity and name of each part they ordered. The first two (+) operators serve the same purposes as in the previous query. The third (+) operator ensures that rows with `NULL` part numbers are also returned:

```

SELECT custname, TO_CHAR(orderdate, 'MON-DD-YYYY') "ORDERDATE",

```

```

quantity, partname
FROM customers, orders, lineitems, parts
WHERE customers.custno = orders.custno (+)
AND orders.orderno = lineitems.orderno (+)
AND lineitems.partno = parts.partno (+);

```

CUSTNAME	ORDERDATE	QUANTITY	PARTNAME
Angelic Co.	OCT-13-1993	15	X-Ray Screen
Angelic Co.	OCT-13-1993	10	Yellow Bag
Angelic Co.	OCT-20-1993	15	X-Ray Screen
Angelic Co.	OCT-27-1993	10	Yellow Bag
Angelic Co.	OCT-27-1993	20	Zoot Suit
Believable Co.	OCT-13-1993	25	X-Ray Screen
Believable Co.	OCT-13-1993	50	Zoot Suit
Believable Co.	OCT-31-1993		
Cables R Us			

**Collection Unnesting Examples** Suppose the database contains a table `HR_INFO` with columns `DEPT`, `LOCATION`, and `MGR`, and a column of nested table type `PEOPLE` which has `NAME`, `DEPT`, and `SAL` columns. You could get all the rows from `HR_INFO` and all the rows from `PEOPLE` using the following statement:

```

SELECT t1.dept, t2.* FROM hr_info t1, TABLE(t1.people) t2
WHERE t2.dept = t1.dept;

```

Now suppose that `PEOPLE` is not a nested table column of `HR_INFO`, but is instead a separate table with columns `NAME`, `DEPT`, `ADDRESS`, `HIREDATE`, and `SAL`. You can extract the same rows as in the preceding example with this statement:

```

SELECT t1.department, t2.*
FROM hr_info t1, TABLE(CAST(MULTISET(
SELECT t3.name, t3.dept, t3.sal FROM people t3
WHERE t3.dept = t1.dept)
AS NESTED_PEOPLE)) t2;

```

Finally, suppose that `PEOPLE` is neither a nested table column of table `HR_INFO` nor a table itself. Instead, you have created a function `PEOPLE_FUNC` that extracts from various sources the name, department, and salary of all employees. You can get the same information as in the preceding examples with the following query:

```

SELECT t1.dept, t2.* FROM HR_INFO t1, TABLE(CAST
(people_func( ... ) AS NESTED_PEOPLE)) t2;

```

**See Also:** *Oracle8i Application Developer's Guide - Fundamentals* for more examples of collection unnesting.

**LEVEL Examples** The following statement returns all employees in hierarchical order. The root row is defined to be the employee whose job is 'PRESIDENT'. The child rows of a parent row are defined to be those who have the employee number of the parent row as their manager number.

```
SELECT LPAD(' ', 2*(LEVEL-1)) || ename org_chart,
       empno, mgr, job
FROM emp
START WITH job = 'PRESIDENT'
CONNECT BY PRIOR empno = mgr;
```

ORG_CHART	EMPNO	MGR	JOB
KING	7839		PRESIDENT
JONES	7566	7839	MANAGER
SCOTT	7788	7566	ANALYST
ADAMS	7876	7788	CLERK
FORD	7902	7566	ANALYST
SMITH	7369	7902	CLERK
BLAKE	7698	7839	MANAGER
ALLEN	7499	7698	SALESMAN
WARD	7521	7698	SALESMAN
MARTIN	7654	7698	SALESMAN
TURNER	7844	7698	SALESMAN
JAMES	7900	7698	CLERK
CLARK	7782	7839	MANAGER
MILLER	7934	7782	CLERK

The following statement is similar to the previous one, except that it does not select employees with the job 'ANALYST'.

```
SELECT LPAD(' ', 2*(LEVEL-1)) || ename org_chart,
       empno, mgr, job
FROM emp
WHERE job != 'ANALYST'
START WITH job = 'PRESIDENT'
CONNECT BY PRIOR empno = mgr;
```

ORG_CHART	EMPNO	MGR	JOB
KING	7839		PRESIDENT
JONES	7566	7839	MANAGER

ADAMS	7876	7788 CLERK
SMITH	7369	7902 CLERK
BLAKE	7698	7839 MANAGER
ALLEN	7499	7698 SALESMAN
WARD	7521	7698 SALESMAN
MARTIN	7654	7698 SALESMAN
TURNER	7844	7698 SALESMAN
JAMES	7900	7698 CLERK
CLARK	7782	7839 MANAGER
MILLER	7934	7782 CLERK

Oracle does not return the analysts SCOTT and FORD, although it does return employees who are managed by SCOTT and FORD.

The following statement is similar to the first one, except that it uses the LEVEL pseudocolumn to select only the first two levels of the management hierarchy:

```
SELECT LPAD(' ',2*(LEVEL-1)) || ename org_chart,
empno, mgr, job
FROM emp
START WITH job = 'PRESIDENT'
CONNECT BY PRIOR empno = mgr AND LEVEL <= 2;
```

ORG_CHART	EMPNO	MGR	JOB
KING	7839		PRESIDENT
JONES	7566	7839	MANAGER
BLAKE	7698	7839	MANAGER
CLARK	7782	7839	MANAGER

**Distributed Query Example** This example shows a query that joins the DEPT table on the local database with the EMP table on the HOUSTON database:

```
SELECT ename, dname
FROM emp@houston, dept
WHERE emp.deptno = dept.deptno;
```

**Correlated Subquery Examples** The following examples show the general syntax of a correlated subquery:

```
SELECT select_list
FROM table1 t_alias1
WHERE expr operator
(SELECT column_list
FROM table2 t_alias2
WHERE t_alias1.column
```

```
        operator t_alias2.column);
UPDATE table1 t_alias1
  SET column =
    (SELECT expr
     FROM table2 t_alias2
     WHERE t_alias1.column = t_alias2.column);
DELETE FROM table1 t_alias1
  WHERE column operator
    (SELECT expr
     FROM table2 t_alias2
     WHERE t_alias1.column = t_alias2.column);
```

The following statement returns data about employees whose salaries exceed their department average. The following statement assigns an alias to EMP, the table containing the salary information, and then uses the alias in a correlated subquery:

```
SELECT deptno, ename, sal
  FROM emp x
  WHERE sal > (SELECT AVG(sal)
               FROM emp
               WHERE x.deptno = deptno)
  ORDER BY deptno;
```

For each row of the EMP table, the parent query uses the correlated subquery to compute the average salary for members of the same department. The correlated subquery performs the following steps for each row of the EMP table:

1. The DEPTNO of the row is determined.
2. The DEPTNO is then used to evaluate the parent query.
3. If that row's salary is greater than the average salary for that row's department, then the row is returned.

The subquery is evaluated once for each row of the EMP table.

**DUAL Table Example** The following statement returns the current date:

```
SELECT SYSDATE FROM DUAL;
```

You could select SYSDATE from the EMP table, but Oracle would return 14 rows of the same SYSDATE, one for every row of the EMP table. Selecting from DUAL is more convenient.

**Sequence Examples** The following statement increments the ZSEQ sequence and returns the new value:




```
SELECT zseq.nextval  
FROM dual;
```

The following statement selects the current value of ZSEQ:

```
SELECT zseq.currval  
FROM dual;
```

## SET CONSTRAINT(S)

### Syntax



### Purpose

To specify, for a particular transaction, whether a deferrable constraint is checked following each DML statement or when the transaction is committed.

### Prerequisites

To specify when a deferrable constraint is checked, you must have `SELECT` privilege on the table to which the constraint is applied unless the table is in your schema.

### Keywords and Parameters

---

<i>constraint</i>	is the name of one or more integrity constraints.
ALL	sets all deferrable constraints for this transaction.
IMMEDIATE	indicates that the conditions specified by the deferrable constraint are checked immediately after each DML statement.
DEFERRED	indicates that the conditions specified by the deferrable constraint are checked when the transaction is committed.

You can verify the success of deferrable constraints prior to committing them by issuing a `SET CONSTRAINTS ALL IMMEDIATE` statement.

---

### Examples

The following statement sets all deferrable constraints in this transaction to be checked immediately following each DML statement:

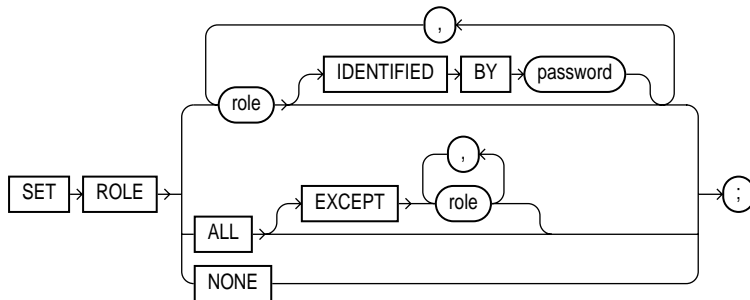
```
SET CONSTRAINTS ALL IMMEDIATE;
```

The following statement checks three deferred constraints when the transaction is committed:

```
SET CONSTRAINTS unq_name, scott.nn_sal,  
                adams.pk_dept@dblink DEFERRED;
```

## SET ROLE

### Syntax



### Purpose

To enable and disable roles for your current session. For information on creating roles, see ["CREATE ROLE"](#) on page 4-365.

When a user logs on, Oracle enables all privileges granted explicitly to the user and all privileges in the user's default roles. During the session, the user or an application can use the `SET ROLE` statement any number of times to change the roles currently enabled for the session. The number of roles that can be concurrently enabled is limited by the initialization parameter `MAX_ENABLED_ROLES`. For information on changing a user's default roles, see ["ALTER USER"](#) on page 7-193

You can see which roles are currently enabled by examining the `SESSION_ROLES` data dictionary view.

### Prerequisites

You must already have been granted the roles that you name in the `SET ROLE` statement.

## Keywords and Parameters

---

<i>role</i>	<p>is a role to be enabled for the current session. Any roles not listed are disabled for the current session.</p> <p><b>Restriction:</b> You cannot specify a role unless it was granted to you either directly or through other roles.</p> <p><code>IDENTIFIED BY password</code> is the password for a role. If the role has a password, you must specify the password to enable the role.</p>
ALL	<p>enables all roles granted to you for the current session except those optionally listed in the EXCEPT clause.</p> <p><b>Restriction:</b> You cannot use this clause to enable roles with passwords that have been granted directly to you.</p> <p><code>EXCEPT</code> Roles listed in the EXCEPT clause must be roles granted directly to you. They cannot be roles granted to you through other roles.</p> <p>If you list a role in the EXCEPT clause that has been granted to you both directly and through another role, the role remains enabled by virtue of the role to which it has been granted.</p>
NONE	<p>disables all roles for the current session, including the DEFAULT role.</p>

---

## Examples

To enable the role GARDENER identified by the password MARIGOLDS for your current session, issue the following statement:

```
SET ROLE gardener IDENTIFIED BY marigolds;
```

To enable all roles granted to you for the current session, issue the following statement:

```
SET ROLE ALL;
```

To enable all roles granted to you except BANKER, issue the following statement:


```
SET ROLE ALL EXCEPT banker;
```

To disable all roles granted to you for the current session, issue the following statement:

```
SET ROLE NONE;
```

## SET TRANSACTION

### Syntax



### Purpose

To establish the current transaction as a read-only or read-write, establish its isolation level, or assign it to a specified rollback segment.

The operations performed by a `SET TRANSACTION` statement affect only your current transaction, not other users or other transactions. Your transaction ends whenever you issue a `COMMIT` or `ROLLBACK` statement. Oracle implicitly commits the current transaction before and after executing a data definition language (DDL) statement.

**See Also:** ["COMMIT"](#) on page 7-230 and ["ROLLBACK"](#) on page 7-564.

### Prerequisites

If you use a `SET TRANSACTION` statement, it must be the first statement in your transaction. However, a transaction need not have a `SET TRANSACTION` statement.

### Keywords and Parameters

#### READ ONLY

establishes the current transaction as a read-only transaction. This clause established **transaction-level read consistency**.

**See Also:** *Oracle8i Concepts*.

All subsequent queries in that transaction only see changes committed before the transaction began. Read-only transactions are useful for reports that run multiple queries against one or more tables while other users update these same tables.

---

	<p><b>Restriction:</b> Only the following statements are permitted in a read-only transaction:</p> <ul style="list-style-type: none"> <li>■ subqueries (that is, <code>SELECT</code> statements without the <i>for_update_clause</i>)</li> <li>■ <code>LOCK TABLE</code></li> <li>■ <code>SET ROLE</code></li> <li>■ <code>ALTER SESSION</code></li> <li>■ <code>ALTER SYSTEM</code></li> </ul>
READ WRITE	<p>establishes the current transaction as a read-write transaction. This clause established <b>statement-level read consistency</b>, which is the default.</p> <p><b>Restriction:</b> You cannot toggle between transaction-level and statement-level read consistency in the same transaction.</p>
ISOLATION LEVEL	<p>specifies how transactions containing database modifications are handled.</p> <p><code>SERIALIZABLE</code> specifies serializable transaction isolation mode as defined in SQL92. If a serializable transaction contains data manipulation language (DML) that attempts to update any resource that may have been updated in a transaction uncommitted at the start of the serializable transaction, then the DML statement fails.</p> <hr/> <p><b>Note:</b> The <code>COMPATIBLE</code> initialization parameter must be set to 7.3.0 or higher for <code>SERIALIZABLE</code> mode to work.</p> <hr/> <p><code>READ COMMITTED</code> is the default Oracle transaction behavior. If the transaction contains DML that requires row locks held by another transaction, then the DML statement waits until the row locks are released.</p>
USE ROLLBACK SEGMENT	<p>assigns the current transaction to the specified rollback segment. This clause also implicitly establishes the transaction as a read-write transaction.</p> <p>This clause lets you to assign transactions of different types to rollback segments of different sizes. For example:</p> <ul style="list-style-type: none"> <li>■ If no long-running queries are concurrently reading the same tables, you can assign small transactions to small rollback segments, which are more likely to remain in memory.</li> <li>■ You can assign transactions that modify tables that are concurrently being read by long-running queries to large rollback segments, so that the rollback information needed for the read-consistent queries is not overwritten.</li> <li>■ You can assign transactions that insert, update, or delete large amounts of data to rollback segments large enough to hold the rollback information for the transaction.</li> </ul> <p>You cannot use the <code>READ ONLY</code> clause and the <code>USE ROLLBACK SEGMENT</code> clause in a single <code>SET TRANSACTION</code> statement or in different statements in the same transaction. Read-only transactions do not generate rollback information and therefore are not assigned rollback segments.</p>

---

## Examples

The following statements could be run at midnight of the last day of every month to count how many ships and containers the company owns. This report would not be affected by any other user who might be adding or removing ships and/or containers.

```
COMMIT;  
SET TRANSACTION READ ONLY;  
SELECT COUNT(*) FROM ship;  
SELECT COUNT(*) FROM container;  
COMMIT;
```

The last `COMMIT` statement does not actually make permanent any changes to the database. It simply ends the read-only transaction.

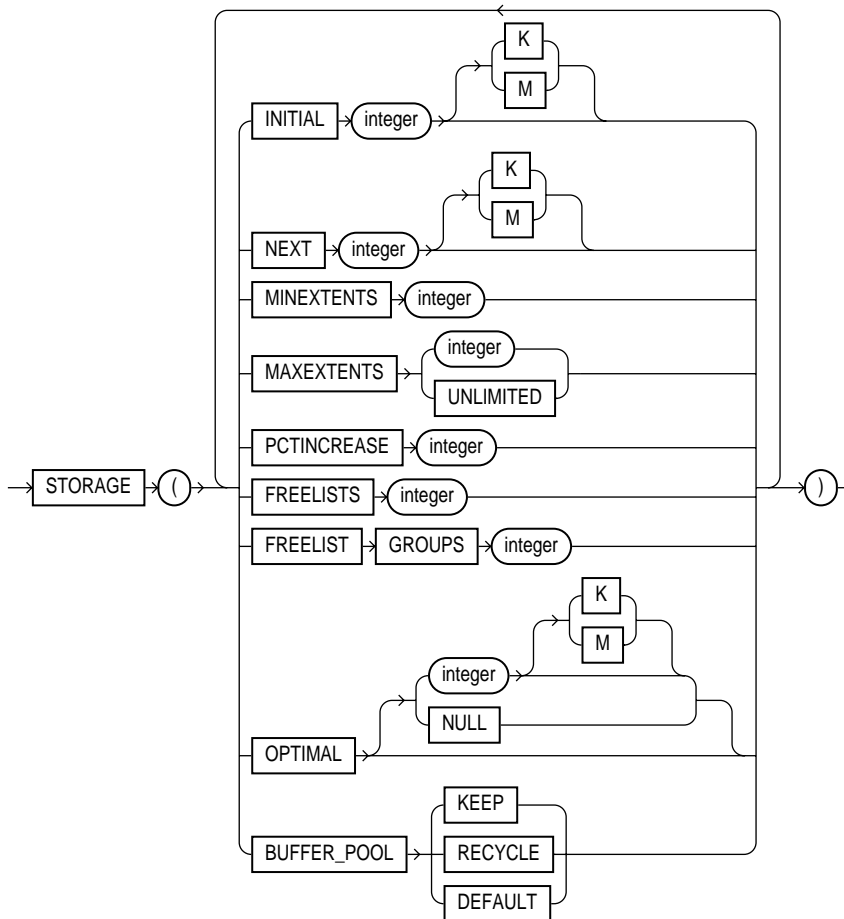
The following statement assigns your current transaction to the rollback segment `OLTP_5`:

```
SET TRANSACTION USE ROLLBACK SEGMENT oltp_5;
```



## storage\_clause

### Syntax



### Purpose

To specify storage characteristics for any of the following schema objects:

- clusters

- indexes
- rollback segments
- materialized views / snapshots
- materialized view logs / snapshot logs
- tables
- tablespaces
- partitions

Storage parameters affect both how long it takes to access data stored in the database and how efficiently space in the database is used. For a discussion of the effects of these parameters, see *Oracle8i Designing and Tuning for Performance*.

When you create a tablespace, you can specify values for the storage parameters. These values serve as default values for segments allocated in the tablespace.

When you alter a tablespace, you can change the values of storage parameters. The new values serve as default values only for subsequently allocated segments (or subsequently created objects).

---

---

**Note:** The *storage\_clause* is interpreted differently for locally managed tablespaces. At creation, Oracle ignores `MAXEXTENTS` and uses the remaining parameter values to calculate the initial size of the segment. For more information, see "[CREATE TABLESPACE](#)" on page 4-419.

---

---

When you create a cluster, index, rollback segment, snapshot, snapshot log, table, or partition, you can specify values for the storage parameters for the segments allocated to these objects. If you omit any storage parameter, Oracle uses the value of that parameter specified for the tablespace.

When you alter a cluster, index, rollback segment, snapshot, snapshot log, table, or partition, you can change the values of storage parameters. The new values affect only future extent allocations.

## Prerequisites

To change the value of a `STORAGE` parameter, you must have the privileges necessary to use the appropriate `CREATE` or `ALTER` statement.

---

## Keywords and Parameters

---

INITIAL	<p>specifies in bytes the size of the object's first extent. Oracle allocates space for this extent when you create the schema object. Use K or M to specify this size in kilobytes or megabytes.</p> <p>The default value is the size of 5 data blocks. The minimum value is the size of 2 data blocks for nonbitmapped segments or 3 data blocks for bitmapped segments, plus one data block for each free list group you specify (see "<a href="#">FREELIST GROUPS</a>" on page 7-608). The maximum value depends on your operating system. Oracle rounds values up to the next multiple of the data block size for values less than 5 data blocks, and rounds up to the next multiple of 5 data blocks for values greater than 5 data blocks.</p> <p><b>Restriction:</b> You cannot specify INITIAL in an ALTER statement.</p>
NEXT	<p>specifies in bytes the size of the next extent to be allocated to the object. Use K or M to specify the size in kilobytes or megabytes. The default value is the size of 5 data blocks. The minimum value is the size of 1 data block. The maximum value depends on your operating system. Oracle rounds values up to the next multiple of the data block size for values less than 5 data blocks. For values greater than 5 data blocks, Oracle rounds up to a value that minimizes fragmentation, as described in <i>Oracle8i Concepts</i>.</p> <p>If you change the value of the NEXT parameter (that is, if you specify it in an ALTER statement), the next allocated extent will have the specified size, regardless of the size of the most recently allocated extent and the value of the PCTINCREASE parameter.</p>
PCTINCREASE	<p>specifies the percent by which the third and subsequent extents grow over the preceding extent. The default value is 50, meaning that each subsequent extent is 50% larger than the preceding extent. The minimum value is 0, meaning all extents after the first are the same size. The maximum value depends on your operating system.</p> <p>Oracle rounds the calculated size of each new extent to the nearest multiple of the data block size.</p> <p>If you change the value of the PCTINCREASE parameter (that is, if you specify it in an ALTER statement), Oracle calculates the size of the next extent using this new value and the size of the most recently allocated extent.</p> <p><b>Suggestion:</b> If you wish to keep all extents the same size, you can prevent SMON from coalescing extents by setting the value of PCTINCREASE to 0. In general, Oracle Corporation recommends a setting of 0 as a way to minimize fragmentation and avoid the possibility of very large temporary segments during processing.</p> <p><b>Restriction:</b> You cannot specify PCTINCREASE for rollback segments. Rollback segments always have a PCTINCREASE value of 0.</p>
MINEXTENTS	<p>specifies the total number of extents to allocate when the object is created. This parameter enables you to allocate a large amount of space when you create an object, even if the space available is not contiguous. The default and minimum value is 1, meaning that Oracle allocates only the initial extent, except for rollback segments, for which the default and minimum value is 2. The maximum value depends on your operating system.</p>

---

If the `MINEXTENTS` value is greater than 1, then Oracle calculates the size of subsequent extents based on the values of the `INITIAL`, `NEXT`, and `PCTINCREASE` parameters.

**Restriction:** You cannot specify `MINEXTENTS` in an `ALTER` statement.

**MAXEXTENTS**

specifies the total number of extents, including the first, that Oracle can allocate for the object. The minimum value is 1 (except for rollback segments, which always have a minimum value of 2). The default value depends on your data block size.

**UNLIMITED**

specifies that extents should be allocated automatically as needed. Oracle Corporation recommends this setting as a way to minimize fragmentation.

However, do not use this clause for rollback segments. Rogue transactions containing inserts, updates, or deletes that continue for a long time will continue to create new extents until a disk is full.

---

**Caution:** A rollback segment that you create without specifying the *storage\_clause* has the same storage parameters as the tablespace that the rollback segment is created in. Thus, if you create the tablespace with `MAXEXTENTS UNLIMITED`, then the rollback segment will also have the same default.

---

**FREELIST  
GROUPS**

specifies the number of groups of free lists for the database object you are creating. The default and minimum value for this parameter is 1. Oracle uses the instance number of Oracle Parallel Server instances to map each instance to one free list group.

Each free list group uses one database block. If you do not specify a large enough value for `INITIAL` to cover the minimum value plus one data block for each free list group, Oracle increases the value of `INITIAL` the necessary amount.

**See Also:** *Oracle8i Parallel Server Concepts*.

**Restriction:** You can specify the `FREELIST GROUPS` parameter only in `CREATE TABLE`, `CREATE CLUSTER`, `CREATE MATERIALIZED VIEW`, `CREATE MATERIALIZED VIEW LOG`, and `CREATE INDEX` statements.

**FREELISTS**

for objects other than tablespaces, specifies the number of free lists for each of the free list groups for the table, partition, cluster, or index. The default and minimum value for this parameter is 1, meaning that each free list group contains one free list. The maximum value of this parameter depends on the data block size. If you specify a `FREELISTS` value that is too large, Oracle returns an error indicating the maximum value.

**Restriction:** You can specify `FREELISTS` in the *storage\_clause* of any statement except when creating or altering a tablespace or rollback segment.

**OPTIMAL**

is relevant only to rollback segments. It specifies an optimal size in bytes for a rollback segment. Use `K` or `M` to specify this size in kilobytes or megabytes. Oracle tries to maintain this size for the rollback segment by dynamically deallocating extents when their data is no longer needed for active transactions. Oracle deallocates as many extents as possible without reducing the total size of the rollback segment below the `OPTIMAL` value.

---

NULL	specifies no optimal size for the rollback segment, meaning that Oracle never deallocates the rollback segment's extents. This is the default behavior.
	The value of <code>OPTIMAL</code> cannot be less than the space initially allocated for the rollback segment specified by the <code>MINEXTENTS</code> , <code>INITIAL</code> , <code>NEXT</code> , and <code>PCTINCREASE</code> parameters. The maximum value depends on your operating system. Oracle rounds values up to the next multiple of the data block size.
BUFFER_POOL	defines a default buffer pool (cache) for a schema object. All blocks for the object are stored in the specified cache. If a buffer pool is defined for a partitioned table or index, then the partitions inherit the buffer pool from the table or index definition, unless overridden by a partition-level definition.
	<b>Note:</b> <code>BUFFER_POOL</code> is not a valid clause for creating or altering tablespaces or rollback segments.
	<b>See Also:</b> <i>Oracle8i Designing and Tuning for Performance</i> for more information about using multiple buffer pools.
KEEP	retains the schema object in memory to avoid I/O operations.
RECYCLE	eliminates blocks from memory as soon as they are no longer needed, thus preventing an object from taking up unnecessary cache space.
DEFAULT	always exists for objects not assigned to <code>KEEP</code> or <code>RECYCLE</code> .

---

## Examples

The following statement creates a table and provides storage parameter values:

```
CREATE TABLE dept
  (deptno    NUMBER(2),
   dname     VARCHAR2(14),
   loc       VARCHAR2(13) )
  STORAGE ( INITIAL 100K NEXT      50K
           MINEXTENTS 1 MAXEXTENTS 50 PCTINCREASE 5);
```

Oracle allocates space for the table based on the `STORAGE` parameter values as follows:

- The `MINEXTENTS` value is 1, so Oracle allocates 1 extent for the table upon creation.
- The `INITIAL` value is 100K, so the first extent's size is 100 kilobytes.
- If the table data grows to exceed the first extent, Oracle allocates a second extent. The `NEXT` value is 50K, so the second extent's size would be 50 kilobytes.

- If the table data subsequently grows to exceed the first two extents, Oracle allocates a third extent. The `PCTINCREASE` value is 5, so the calculated size of the third extent is 5% larger than the second extent, or 52.5 kilobytes. If the data block size is 2 kilobytes, Oracle rounds this value to 52 kilobytes.

If the table data continues to grow, Oracle allocates more extents, each 5% larger than the previous one.

- The `MAXEXTENTS` value is 50, so Oracle can allocate as many as 50 extents for the table.

The following statement creates a rollback segment and provides storage parameter values:

```
CREATE ROLLBACK SEGMENT rsone
  STORAGE ( INITIAL 10K NEXT 10K
           MINEXTENTS 2 MAXEXTENTS 25
           OPTIMAL 50K );
```

Oracle allocates space for the rollback segment based on the `STORAGE` parameter values as follows:

- The `MINEXTENTS` value is 2, so Oracle allocates 2 extents for the rollback segment upon creation.
- The `INITIAL` value is 10K, so the first extent's size is 10 kilobytes.
- The `NEXT` value is 10K, so the second extent's size is 10 kilobytes.
- If the rollback data exceeds the first two extents, Oracle allocates a third extent. The `PCTINCREASE` value for rollback segments is always 0, so the third and subsequent extents are the same size as the second extent, 10 kilobytes.
- The `MAXEXTENTS` value is 25, so Oracle can allocate as many as 25 extents for the rollback segment.
- The `OPTIMAL` value is 50K, so Oracle deallocates extents if the rollback segment exceeds 50 kilobytes. Oracle deallocates only extents that contain data for transactions that are no longer active.

# TRUNCATE

---



---

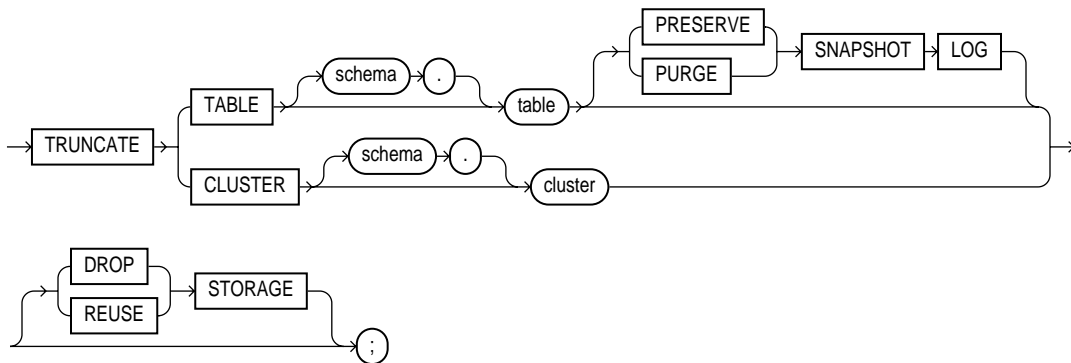
**WARNING:** You cannot roll back a TRUNCATE statement.

---



---

## Syntax



## Purpose

To remove all rows from a table or cluster and reset the `STORAGE` parameters to the values when the table or cluster was created.

Deleting rows with the `TRUNCATE` statement can be more efficient than dropping and re-creating a table. Dropping and re-creating a table invalidates the table's dependent objects, requires you to regrant object privileges on the table, and requires you to re-create the table's indexes, integrity constraint, and triggers and respecify its storage parameters. Truncating has none of these effects.

**See Also:** ["DELETE"](#) on page 7-464, ["DROP CLUSTER"](#) on page 7-472, and ["DROP TABLE"](#) on page 7-501.

## Prerequisites

The table or cluster must be in your schema or you must have `DROP ANY TABLE` system privilege.

## Keywords and Parameters

---

TABLE

specifies the schema and name of the table to be truncated. This table cannot be part of a cluster. If you omit *schema*, Oracle assumes the table is in your own cluster.

You can truncate index-organized tables and temporary tables. When you truncate a temporary table, only the rows created during the current session are truncated.

The table's storage parameter `NEXT` is changed to be the size of the last extent deleted from the segment in the process of truncation.

Oracle also automatically truncates and resets any existing `UNUSABLE` indicators for the following indexes on *table*: range and hash partitions of local indexes and subpartitions of local indexes.

If *table* is not empty, Oracle marks `UNUSABLE` all nonpartitioned indexes and all partitions of global partitioned indexes on the table.

For a domain index, this statement invokes the appropriate truncate routine to truncate the domain index data.

**See Also:** *Oracle8i Data Cartridge Developer's Guide*.

If *table* (whether it is a regular or index-organized table) contains `LOB` columns, all `LOB` data and `LOB` index segments will be truncated.

If *table* is partitioned, all partitions or subpartitions, as well as the `LOB` data and `LOB` index segments for each partition or subpartition, will be truncated.

---

**Note:** When you truncate a table, Oracle automatically deletes all data in the table's indexes and any materialized view direct-load `INSERT` information held in association with the table. (This information is independent of any materialized view/snapshot log.) If this direct-load `INSERT` information is deleted, an incremental refresh of the materialized view may lose data.

---

### Restrictions:

- You cannot individually truncate a table that is part of a cluster. You must either truncate the cluster, delete all rows from the table, or drop and re-create the table.
- You cannot truncate the parent table of an enabled referential integrity constraint. You must disable the constraint before truncating the table. (An exception is that you may truncate the table if the integrity constraint is self-referential.)
- You cannot truncate a table if any domain indexes defined on any of its columns are marked `LOADING` or `FAILED`.



---

SNAPSHOT LOG	<p>specifies whether a snapshot log defined on the table is to be preserved or purged when the table is truncated. This clause allows snapshot master tables to be reorganized through export/import without affecting the ability of primary-key snapshots defined on the master to be fast refreshed. To support continued fast refresh of primary-key snapshots, the snapshot log must record primary-key information.</p> <p><b>See Also:</b> <i>Oracle8i Replication</i> for more information about snapshot logs and the TRUNCATE statement.</p>
PRESERVE	specifies that any snapshot log should be preserved when the master table is truncated. This is the default.
PURGE	specifies that any snapshot log should be purged when the master table is truncated.
CLUSTER	<p>specifies the schema and name of the cluster to be truncated. You can truncate only an indexed cluster, not a hash cluster. If you omit <i>schema</i>, Oracle assumes the table is in your own cluster.</p> <p>When you truncate a cluster, Oracle also automatically deletes all data in the cluster's tables' indexes.</p>
DROP STORAGE	deallocates all space from the deleted rows from the table or cluster except the space allocated by the table's or cluster's MINEXTENTS parameter. This space can subsequently be used by other objects in the tablespace. This is the default.
REUSE STORAGE	<p>retains the space from the deleted rows allocated to the table or cluster. Storage values are not reset to the values when the table or cluster was created. This space can subsequently be used only by new data in the table or cluster resulting from inserts or updates.</p> <p>The DROP STORAGE clause and REUSE STORAGE clause also apply to the space freed by the data deleted from associated indexes.</p>

---

**Note:** If you have specified more than one free list for the object you are truncating, the REUSE STORAGE clause also removes any mapping of free lists to instances, and resets the high-water mark to the beginning of the first extent.

---

## Examples

The following statement deletes all rows from the EMP table and returns the freed space to the tablespace containing EMP:

```
TRUNCATE TABLE emp;
```

The above statement also deletes all data from all indexes on EMP and returns the freed space to the tablespaces containing them.

The following statement deletes all rows from all tables in the CUST cluster, but leaves the freed space allocated to the tables:

```
TRUNCATE CLUSTER cust REUSE STORAGE
```


The above statement also deletes all data from all indexes on the tables in `CUST`.

The following statements are examples of truncate statements that preserve snapshot logs:

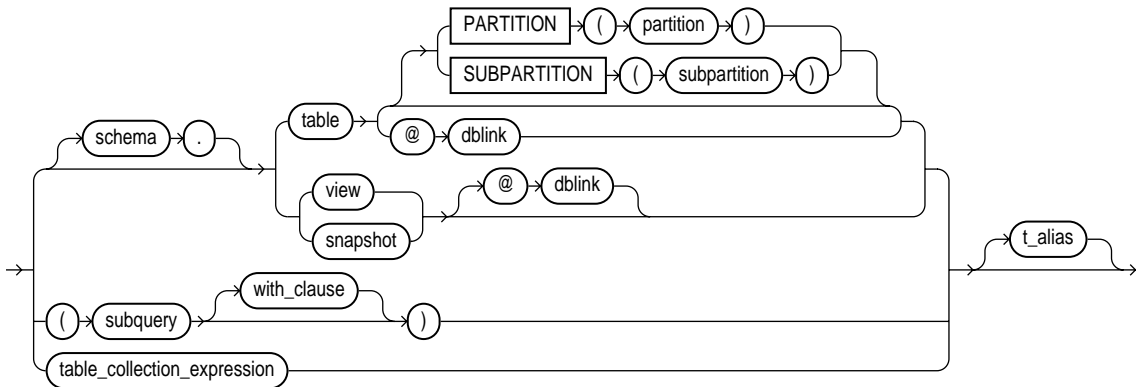
```
TRUNCATE TABLE emp PRESERVE SNAPSHOT LOG;  
TRUNCATE TABLE stock;
```

# UPDATE

## Syntax




### DML\_table\_expression\_clause::=




**subquery:** see "SELECT and Subqueries" on page 7-569.

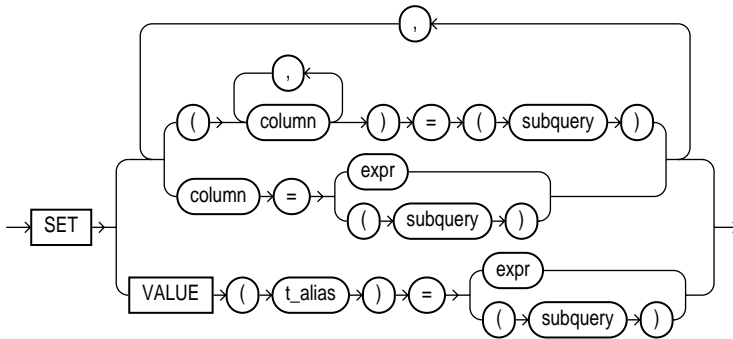
### with\_clause::=



### table\_collection\_expression::=




**set\_clause::=**



**where\_clause::=**



**returning\_clause::=**



**Purpose**

To change existing values in a table or in a view's base table.

**Prerequisites**

For you to update values in a table, the table must be in your own schema or you must have `UPDATE` privilege on the table.

For you to update values in the base table of a view,

- You must have `UPDATE` privilege on the view, and
- Whoever owns the schema containing the view must have `UPDATE` privilege on the base table.

If the `SQL92_SECURITY` initialization parameter is set to `TRUE`, then you must have `SELECT` privilege on the table whose column values you are referencing (such as the columns in a *where\_clause*) to perform an `UPDATE`.

---

The UPDATE ANY TABLE system privilege also allows you to update values in any table or any view's base table.

## Keywords and Parameters

---

<i>hint</i>	<p>is a comment that passes instructions to the optimizer on choosing an execution plan for the statement.</p> <p>You can place a parallel hint immediately after the UPDATE keyword to parallelize both the underlying scan and UPDATE operations.</p> <p><b>See Also:</b></p> <ul style="list-style-type: none"><li>■ <i>Oracle8i Designing and Tuning for Performance</i> and "<a href="#">Hints</a>" on page 2-63 for the syntax and description of hints.</li><li>■ <i>Oracle8i Designing and Tuning for Performance</i>, <i>Oracle8i Parallel Server Concepts</i>, and <i>Oracle8i Concepts</i> for detailed information about parallel DML.</li></ul>
<i>DML_table_expression_clause</i>	
<i>schema</i>	<p>is the schema containing the table or view. If you omit <i>schema</i>, Oracle assumes the table or view is in your own schema.</p>
<i>table   view   subquery</i>	<p>is the name of the table or view, or the columns returned by a subquery, to be updated. Issuing an UPDATE statement against a table fires any UPDATE triggers associated with the table. If you specify <i>view</i>, Oracle updates the view's base table.</p> <p>If <i>table</i> (or the base table of <i>view</i>) contains one or more domain index columns, this statement executes the appropriate indextype update routine.</p> <p><b>See Also:</b> <i>Oracle8i Data Cartridge Developer's Guide</i> for more information on these routines.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ You cannot execute this statement if <i>table</i> (or the base table of <i>view</i>) contains any domain indexes marked LOADING or FAILED.</li><li>■ You cannot specify the <i>order_by_clause</i> in the subquery of the <i>DML_query_expression_clause</i>.</li></ul>

- You cannot update a view except with `INSTEAD OF` triggers if the view's defining query contains one of the following constructs:
  - A set operator
  - A `DISTINCT` operator
  - An aggregate function
  - A `GROUP BY`, `ORDER BY`, `CONNECT BY`, or `START WITH` clause
  - A collection expression in a `SELECT` list
  - A subquery in a `SELECT` list
  - Joins (with some exceptions). See *Oracle8i Administrator's Guide* for details.
- If a view was created with the `WITH CHECK OPTION`, you can update the view only if the resulting data satisfies the view's defining query.
- If you specify an index, index partition, or index subpartition that has been marked `UNUSABLE`, the `UPDATE` statement will fail unless the `SKIP_UNUSABLE_INDEXES` parameter has been set to `TRUE`.

**See Also:** "[ALTER SESSION](#)" on page 7-83

`PARTITION` (*partition*) / `SUBPARTITION` (*subpartition*)

specifies the name of the partition or subpartition within *table* targeted for updates. You need not specify the partition name when updating values in a partitioned table. However in some cases specifying the partition name can be more efficient than a complicated *where\_clause*.

*dblink*

is a complete or partial name of a database link to a remote database where the table or view is located. For information on referring to database links, see "[Referring to Objects in Remote Databases](#)" on page 2-79. You can use a database link to update a remote table or view only if you are using Oracle's distributed functionality.

If you omit *dblink*, Oracle assumes the table or view is on the local database.

*with\_clause*

restricts the subquery in one of the following ways:

- `WITH READ ONLY` specifies that the subquery cannot be updated.
- `WITH CHECK OPTION` specifies that Oracle prohibits any changes to that table that would produce rows that are not included in the subquery.

**See Also:** [WITH CHECK OPTION Example](#) on page 7-587.

*table\_collection\_expression*

informs Oracle that the collection value expression should be treated as a table. You can use a *table\_collection\_expression* to update rows in one table based on rows from another table. For example, you could roll up four quarterly sales tables into a yearly sales table.

*collection\_expression* is a subquery that selects a nested table column from *table* or *view*.

---

**Note:** In earlier releases of Oracle, *table\_collection\_expr* was expressed as "THE subquery". That usage is now deprecated.

---

<i>t_alias</i>	provides a <b>correlation name</b> for the table, view, or subquery to be referenced elsewhere in the statement.
	<b>Note:</b> This alias is <b>required</b> if the <i>DML_query_expression_clause</i> references any object type attributes or object type methods.
<i>set_clause</i>	<p><i>column</i> is the name of a column of the table or view that is to be updated. If you omit a column of the table from the <i>set_clause</i>, that column's value remains unchanged.</p> <p><b>Restrictions:</b></p> <ul style="list-style-type: none"><li>■ If <i>column</i> refers to a LOB object attribute, you must first initialize it with a value of empty or null. You cannot update it with a literal. Also, if you are updating a LOB value using some method other than a direct UPDATE SQL statement, you must first lock the row containing the LOB.</li></ul> <p><b>See Also:</b> <a href="#">LOB Locking Example</a> on page 7-585.</p> <ul style="list-style-type: none"><li>■ If <i>column</i> is part of the partitioning key of a partitioned table, UPDATE will fail if you change a value in the column that would move the row to a different partition or subpartition, unless you enable row movement.</li></ul> <p><b>See Also:</b> The <i>row_movement_clause</i> of "<a href="#">CREATE TABLE</a>" on page 4-381 or "<a href="#">ALTER TABLE</a>" on page 7-123.</p>
<i>subquery</i>	<p>is a subquery that returns exactly one row for each row updated.</p> <ul style="list-style-type: none"><li>■ If you specify only one column in the <i>set_clause</i>, the subquery can return only one value.</li><li>■ If you specify multiple columns in the <i>set_clause</i>, the subquery must return as many values as you have specified columns.</li></ul> <p>If the subquery returns no rows, then the column is assigned a null.</p> <p><b>See Also:</b> "<a href="#">SELECT and Subqueries</a>" on page 7-569 and "<a href="#">Using Subqueries</a>" on page 5-24.</p> <hr/> <p><b>Note:</b> If this <i>subquery</i> refers to remote objects, the UPDATE operation can run in parallel as long as the reference does not loop back to an object on the local database. (See the description of the <i>parallel_clause</i> in "<a href="#">CREATE TABLE</a>" on page 4-381.) However, if the <i>subquery</i> in the <i>DML_query_expression_clause</i> refers to any remote objects, the UPDATE operation will run serially without notification.</p>
<i>expr</i>	<p>is the new value assigned to the corresponding column. This expression can contain host variables and optional indicator variables.</p> <p><b>See Also:</b> The syntax description in "<a href="#">Expressions</a>" on page 5-1.</p>

---

VALUE	lets you specify the entire row of an object table. <b>Restriction:</b> You can specify this clause only for an object table. <b>See Also:</b> "SET VALUE Example" on page 7-622.
-------	---

---

**Note:** If you insert string literals into a RAW column, during subsequent queries, Oracle will perform a full table scan rather than using any index that might exist on the RAW column.

---

<i>where_clause</i>	restricts the rows updated to those for which the specified <i>condition</i> is TRUE. If you omit this clause, Oracle updates all rows in the table or view.  The <i>where_clause</i> determines the rows in which values are updated. If you do not specify the <i>where_clause</i> , all rows are updated. For each row that satisfies the <i>where_clause</i> , the columns to the left of the equals (=) operator in the <i>set_clause</i> are set to the values of the corresponding expressions on the right. The expressions are evaluated as the row is updated. <b>See Also:</b> The syntax description of "Conditions" on page 5-14.
---------------------	---

---

<i>returning_clause</i>	retrieves the rows affected by the UPDATE statement. <ul style="list-style-type: none"> <li>■ When you are updating a single row, this clause can retrieve column expressions that use the updated columns of the row, rowid, and REFS to the updated row and store them in PL/SQL variables or bind variables.</li> <li>■ When you are updating multiple rows, this clause can store the values from expressions, rowid, and REFS involving the updated rows in bind arrays.</li> <li>■ You can also use UPDATE with a <i>returning_clause</i> to update from views with single base tables.</li> </ul>
-------------------------	--

---

<i>expr list</i>	is some of the syntax descriptions in "Expressions" on page 5-1. You must specify a column expression in the <i>expr list</i> for each variable in the <i>data_item list</i> .
------------------	--

---

INTO	indicates that the values of the changed rows are to be stored in the <i>data_item</i> variable(s) specified in <i>data_item list</i> .
------	---

---

<i>data_item</i>	is a PL/SQL variable or bind variable which stores the retrieved <i>expr</i> value in the <i>expr list</i> .
------------------	--

**Restrictions:**

- You cannot use this clause with parallel DML or with remote objects.
- You cannot retrieve LONG types with this clause.
- You cannot specify this clause for a view on which an INSTEAD OF trigger has been defined.

---

## Examples

**Simple Examples** The following statement gives null commissions to all employees with the job TRAINEE:



```
UPDATE emp
  SET comm = NULL
  WHERE job = 'TRAINEE';
```

The following statement promotes JONES to manager of Department 20 with a \$1,000 raise (assuming there is only one JONES):

```
UPDATE emp
  SET job = 'MANAGER', sal = sal + 1000, deptno = 20
  WHERE ename = 'JONES';
```

The following statement increases the balance of bank account number 5001 in the ACCOUNTS table on a remote database accessible through the database link BOSTON:

```
UPDATE accounts@boston
  SET balance = balance + 500
  WHERE acc_no = 5001;
```

**PARTITION Example** The following example updates values in a single partition of the SALES table:

```
UPDATE sales PARTITION (feb96) s
  SET s.account_name = UPPER(s.account_name);
```

**Complex Example** This example shows the following syntactic constructs of the UPDATE statement:

- Both forms of the *set\_clause* together in a single statement
- A correlated subquery
- A *where\_clause* to limit the updated rows

```
UPDATE emp a
  SET deptno =
    (SELECT deptno
     FROM dept
     WHERE loc = 'BOSTON'),
    (sal, comm) =
    (SELECT 1.1*AVG(sal), 1.5*AVG(comm)
     FROM emp b
     WHERE a.deptno = b.deptno)
  WHERE deptno IN
    (SELECT deptno
     FROM dept
     WHERE loc = 'DALLAS'
      OR loc = 'DETROIT');
```

The above UPDATE statement performs the following operations:

- Updates only those employees who work in Dallas or Detroit
- Sets DEPTNO for these employees to the DEPTNO of Boston
- Sets each employee's salary to 1.1 times the average salary of their department
- Sets each employee's commission to 1.5 times the average commission of their department

**SET VALUE Example** The following statement updates a row of object table TABLE1 by selecting a row from another object table TABLE2:

```
UPDATE table1 p SET VALUE(p) =
    SELECT VALUE(q) FROM table2 q WHERE p.id = q.id)
WHERE p.id = 10;
```

The subquery uses the VALUE object reference function in its expression.

**Correlated Update Example** The following example updates particular rows of the PROJS nested table corresponding to the department whose department equals 123:

```
UPDATE TABLE(SELECT projs
    FROM dept d WHERE d.dno = 123) p
SET p.budgets = p.budgets + 1
WHERE p.pno IN (123, 456);
```

**RETURNING Example** The following example returns values from the updated row and stores the result in PL/SQL variables BND1, BND2, BND3:

```
UPDATE emp
SET job = 'MANAGER', sal = sal + 1000, deptno = 20
WHERE ename = 'JONES'
RETURNING sal*0.25, ename, deptno INTO bnd1, bnd2, bnd3;
```

---

---

# Syntax Diagrams

*One picture is worth a thousand words.*

Anonymous

*Syntax diagrams* are drawings that illustrate valid SQL syntax. To read a diagram, trace it from left to right, in the direction shown by the arrows.


Commands and other keywords appear in UPPERCASE inside rectangles. Type them exactly as shown in the rectangles. Parameters appear in lowercase inside ovals. Variables are used for the parameters. Punctuation, operators, delimiters, and terminators appear inside circles.

If the syntax diagram has more than one path, you can choose any path to travel.

If you have the choice of more than one keyword, operator, or parameter, your options appear in a vertical list.

## Required Keywords and Parameters

Required keywords and parameters can appear singly or in a vertical list of alternatives. Single required keywords and parameters appear on the *main path*, that is, on the horizontal line you are currently traveling. In the following example, *library\_name* is a required parameter:




If there is a library named HQ\_LIB, then, according to the diagram, the following statement is valid:

```
DROP LIBRARY hq_lib;
```


---

If multiple keywords or parameters appear in a vertical list that intersects the main path, one of them is required. That is, you must choose one of the keywords or parameters, but not necessarily the one that appears on the main path. In the following example, you must choose one of the four settings:



## Optional Keywords and Parameters

If keywords and parameters appear in a vertical list *above* the main path, they are optional. In the following example, instead of traveling down a vertical line, you can continue along the main path:




According to the diagram, all of the following statements are valid:

```
DEALLOCATE UNUSED;  
DEALLOCATE UNUSED KEEP 1000;  
DEALLOCATE UNUSED KEEP 10M;
```

## Syntax Loops


Loops let you repeat the syntax within them as many times as you like. In the following example, after choosing one expression, you can go back repeatedly to choose another, separated by commas.



---

## Multipart Diagrams

Read a multipart diagram as if all the main paths were joined end to end. The following example is a two-part diagram:



According to the diagram, the following statement is valid:

```
CREATE OUTLINE ON UPDATE;
```

## Database Objects

The names of Oracle identifiers, such as tables and columns, must not exceed 30 characters in length. The first character must be a letter, but the rest can be any combination of letters, numerals, dollar signs (\$), pound signs (#), and underscores (\_).

However, if an Oracle identifier is enclosed by double quotation marks ("), it can contain any combination of legal characters, including spaces but excluding quotation marks.

Oracle identifiers are not case-sensitive except when enclosed by double quotation marks.

For more information, see ["Schema Object Naming Rules"](#) on page 2-72.



---

---

## Oracle and Standard SQL

*High thoughts must have high language.*

Aristophanes, *Frogs*

This appendix discusses Oracle's conformance to the SQL-92 standards established by industry standards governing bodies. We are assessing the changes that will be required to this appendix in light of the new SQL-99 standards. The appendix also described how to locate extensions to standard SQL with the FIPS Flagger.

### Conformance with Standard SQL

This section declares Oracle's conformance to the SQL standards established by these organizations:

- American National Standards Institute (ANSI)
- International Standards Organization (ISO)
- United States Federal Government

### ANSI and ISO Compliance

Oracle8i complies at the Entry level as defined in the ANSI document, X3.135-1992, "Database Language SQL." You can obtain a copy of the ANSI standard from this address:

American National Standards Institute  
1430 Broadway  
New York, NY 10018 USA

The ANSI and ISO SQL standards require conformance claims to state the type of conformance and the implemented facilities. The Oracle server, Oracle Precompilers

for C/C++ Release 8.1, Oracle Precompiler for Cobol Release 8.1, and SQL\*Module for ADA Release 8.0.4 provide conformance with the ANSI X3.135-1992/ISO 9075-1992 standard:

- Compliance at Entry Level (including both SQL-DDL and SQL-DML)
- Module Language for ADA
- Embedded SQL C
- Embedded SQL COBOL

In addition to full compliance at the Entry level, Oracle complies partially at the Transitional, Intermediate, and Full levels as described in [Table B-1](#) (including both SQL-DDL and SQL-DML).

**Table B-1 Oracle Compliance at Transitional, Intermediate, and Full Levels**

Level	SQL92 Feature (number and name)
<b>Transitional</b>	7. TRIM function
	8. UNION in views
	9. Implicit numeric casting
	10. Implicit character casting
	13. Grouped operations
	14. Qualified * in SELECT list
	15. Lowercase identifiers
	16. PRIMARY KEY enhancement
	18. Multiple module support
<b>Intermediate</b>	21. INSERT expressions
	31. Schema definition statement
	42. National character
<b>Full</b>	48. Expanded null predicate
	60. Trailing underscore
	62. Referential name order



## FIPS Compliance

Oracle complies completely with FIPS PUB 127-2 for Entry SQL. In addition, the following information is provided for Section 16, "Special Procurement Considerations."

### Section 16.2 Programming Language Interfaces

The Oracle precompilers support the use of embedded SQL in C and COBOL. SQL\*Module supports the use of Module Language in ADA.

### Section 16.3 Style of Language Interface

Oracle with SQL\*Module supports Module Language for Ada. Oracle with the Oracle precompilers supports C and COBOL. The specific languages supported depend on your operating system.

### Section 16.5 Interactive Direct SQL

Oracle8i with SQL\*Plus Version 3.1 (as well as other Oracle tools) supports "direct invocation" of the following SQL statements, meeting the requirements of FIPS PUB 127-2:

- CREATE TABLE statement
- CREATE VIEW statement
- GRANT statement
- INSERT statement
- SELECT statement, with ORDER BY clause but not INTO clause
- UPDATE statement: searched
- DELETE statement: searched
- COMMIT WORK statement
- ROLLBACK WORK statement

Most other SQL statements described in this reference are also supported interactively.

## Section 16.6 Sizing for Database Constructs

Table B-2 lists requirements identified in FIPS PUB 127-1 and how they are met by Oracle8i.

**Table B-2** *Sizing for Database Constructs*

Database Constructs	FIPS	Oracle8i
Length of an identifier (in bytes)	18	30
Length of CHARACTER datatype (in bytes)	240	2000
Decimal precision of NUMERIC datatype	15	38
Decimal precision of DECIMAL datatype	15	38
Decimal precision of INTEGER datatype	9	38
Decimal precision of SMALLINT datatype	4	38
Binary precision of FLOAT datatype	20	126
Binary precision of REAL datatype	20	63
Binary precision of DOUBLE PRECISION datatype	30	126
Columns in a table	100	1000
Values in an INSERT statement	100	1000
SET clauses in an UPDATE statement <sup>(a)</sup>	20	1000
Length of a row <sup>(b,c)</sup>	2,000	2,000,000
Columns in a UNIQUE constraint	6	32
Length of a UNIQUE constraint <sup>(b)</sup>	120	(d)
Length of foreign key column list <sup>(b)</sup>	120	(d)
Columns in a GROUP BY clause	6	255 <sup>(e)</sup>
Length of GROUP BY column list	120	(e)
Sort specifications in ORDER BY clause	6	255 <sup>(e)</sup>
Length of ORDER BY column list	120	(e)
Columns in a referential integrity constraint	6	32
Tables referenced in a SQL statement	15	No limit
Cursors simultaneously open	10	(f)
Items in a SELECT list	100	1000

**Table B-2 (Cont.) Sizing for Database Constructs**

- 
- (a) The number of `SET` clauses in an `UPDATE` statement refers to the number items separated by commas following the `SET` keyword.
  - (b) The FIPS PUB defines the length of a collection of columns to be the sum of: twice the number of columns, the length of each character column in bytes, decimal precision plus 1 of each exact numeric column, binary precision divided by 4 plus 1 of each approximate numeric column.
  - (c) The Oracle limit for the maximum row length is based on the maximum length of a row containing a `LONG` value of length 2 gigabytes and 999 `VARCHAR2` values, each of length 4000 bytes:  $2(254) + 231 + (999(4000))$ .
  - (d) The Oracle limit for a `UNIQUE` key is half the size of an Oracle data block (specified by the initialization parameter `DB_BLOCK_SIZE`) minus some overhead.
  - (e) Oracle places no limit on the number of columns in a `GROUP BY` clause or the number of sort specifications in an `ORDER BY` clause. However, the sum of the sizes of all the expressions in either a `GROUP BY` clause or an `ORDER BY` clause is limited to the size of an Oracle data block (specified by the initialization parameter `DB_BLOCK_SIZE`) minus some overhead.
  - (f) The Oracle limit for the number of cursors simultaneously opened is specified by the initialization parameter `OPEN_CURSORS`. The maximum value of this parameter depends on the memory available on your operating system and exceeds 100 in all cases.
- 

## Section 16.7 Character Set Support

Oracle supports the ASCII character set (FIPS PUB 1-2) on most computers and the EBCDIC character set on IBM mainframe computers. Oracle supports both single-byte and multibyte character sets.

## Oracle Extensions to Standard SQL

Oracle supports numerous features that extend beyond standard SQL. In your Oracle applications, you can use these extensions just as you can use Entry SQL92.

If you are concerned with the portability of your applications to other implementations of SQL, use Oracle's FIPS Flagger to locate Oracle extensions to Entry SQL92 in your embedded SQL programs. The FIPS Flagger is part of the Oracle precompilers and the SQL\*Module compiler.

**See Also:** *Pro\*COBOL Precompiler Programmer's Guide* and *Pro\*C/C++ Precompiler Programmer's Guide* for information on how to use the FIPS Flagger.



---



---

## Oracle Reserved Words

*The words I use are everyday words and yet are not the same!*

Paul Claudel, *La Muse Qui Est la Grace*

This appendix lists Oracle reserved words. Words followed by an asterisk (\*) are also ANSI reserved words.

---



---

**Note:** In addition to the following reserved words, Oracle uses system-generated names beginning with "SYS\_" for implicitly generated schema objects and subobjects. Oracle discourages you from using this prefix in the names you explicitly provide to your schema objects and subobjects to avoid possible conflict in name resolution.

---



---

**Table C-1 Oracle Reserved Words**

---

ACCESS	CHAR *	DEFAULT *
ADD *	CHECK *	DELETE *
ALL *	CLUSTER	DESC *
ALTER *	COLUMN	DISTINCT *
AND *	COMMENT	DROP *
ANY *	COMPRESS	ELSE *
AS *	CONNECT *	EXCLUSIVE
ASC *	CREATE *	EXISTS
AUDIT	CURRENT *	FILE
BETWEEN *	DATE *	FLOAT *
BY *	DECIMAL *	FOR *

---

**Table C-1 Oracle Reserved Words**

---

FROM *	NOT *	SHARE
GRANT *	NOWAIT	SIZE *
GROUP *	NULL *	SMALLINT *
HAVING *	NUMBER	START
IDENTIFIED	OF *	SUCCESSFUL
IMMEDIATE *	OFFLINE	SYNONYM
IN *	ON *	SYSDATE
INCREMENT	ONLINE	TABLE *
INDEX	OPTION *	THEN *
INITIAL	OR *	TO *
INSERT *	ORDER *	TRIGGER
INTEGER *	PCTFREE	UID
INTERSECT *	PRIOR *	UNION *
INTO *	PRIVILEGES *	UNIQUE *
IS *	PUBLIC *	UPDATE *
LEVEL *	RAW	USER *
LIKE *	RENAME	VALIDATE
LOCK	RESOURCE	VALUES *
LONG	REVOKE *	VARCHAR *
MAXEXTENTS	ROW	VARCHAR2
MINUS	ROWID	VIEW *
MLSLABEL	ROWNUM	WHENEVER *
MODE	ROWS *	WHERE
MODIFY	SELECT *	WITH *
NOAUDIT	SESSION *	
NOCOMPRESS	SET *	

---

## Symbols

---

- \$ number format element, 2-41
- (+) operator, 3-16
- , (comma)
  - date format element, 2-46
  - number format element, 2-41
- : (colon) date format element, 2-46
- (dash) date format element, 2-46
- ; (semicolon) date format element, 2-46
- / (slash) date format element, 2-46
- (period)
  - date format element, 2-46
  - number format element, 2-41

## Numerics

---

- 0 number format element, 2-41
- 20th century, 2-47, 2-49
  - specifying, 2-49
- 21st century, 2-47, 2-49
  - specifying, 2-49
- 8 number format element, 2-41
- 9 number format element, 2-41

## A

---

- ABS function, 4-12
- ABSI
  - standards, B - 1
- ACCOUNT LOCK clause
  - of ALTER USER. *See* CREATE USER
  - of CREATE USER, 7-453
- ACCOUNT UNLOCK clause
  - of ALTER USER. *See* CREATE USER
  - of ALTER USER. *See* CREATE USER.
  - of CREATE USER, 7-453
- ACOS function, 4-13
- ACTIVATE STANDBY DATABASE clause
  - of ALTER DATABASE, 7-11
- AD (A.D.) date format element, 2-46, 2-48
- ADD clause
  - of ALTER DIMENSION, 7-27
  - of ALTER TABLE, 7-140
- ADD DATAFILE clause
  - of ALTER TABLESPACE, 7-179
- ADD LOGFILE clause
  - of ALTER DATABASE, 7-9
- ADD LOGFILE GROUP clause
  - of ALTER DATABASE, 7-18
- ADD LOGFILE MEMBER clause
  - of ALTER DATABASE, 7-9, 7-18
- ADD LOGFILE THREAD clause
  - of ALTER DATABASE, 7-18
- ADD OVERFLOW clause
  - of ALTER TABLE, 7-155
- ADD PARTITION, 7-160
- ADD PARTITION clause
  - of ALTER TABLE, 7-158, 7-160
- ADD PRIMARY KEY clause
  - of ALTER MATERIALIZED VIEW LOG, 7-61
- ADD ROWID clause
  - of ALTER MATERIALIZED VIEW, 7-61
  - of ALTER MATERIALIZED VIEW LOG, 7-61
- ADD TEMPFILE clause
  - of ALTER TABLESPACE, 7-179
- ADD\_MONTHS function, 4-13
- ADMINISTER ANY TRIGGER system

- privilege, 7-527
- ADVISE clause
  - of ALTER SESSION, 7-84
- AFTER clause
  - of CREATE TRIGGER, 7-428
- AFTER triggers, 7-428
- aggregate functions, 4-6
- aliases
  - for columns, 5-19
  - for expressions in view query, 7-458
  - specifying in queries and subqueries, 7-576
- ALL clause
  - of SELECT, 7-573
  - of SET CONSTRAINTS, 7-598
  - of SET ROLE, 7-601
- ALL EXCEPT clause
  - of SET ROLE, 7-601
- ALL operator, 3-6
- ALL PRIVILEGES clause
  - of GRANT object\_privileges, 7-533
  - of REVOKE schema\_object\_privileges, 7-560
- ALL PRIVILEGES shortcut
  - of AUDIT sql\_statements, 7-214
- ALL shortcut
  - of AUDIT sql\_statements, 7-214
- ALL\_COL\_COMMENTS view, 7-228
- ALL\_ROWS hint, 2-64
- ALL\_TAB\_COMMENTS view, 7-228
- ALLOCATE EXTENT clause
  - of ALTER CLUSTER, 7-3, 7-4
  - of ALTER INDEX, 7-31, 7-35
  - of ALTER TABLE, 7-150
- ALTER ANY CLUSTER system privilege, 7-521
- ALTER ANY DIMENSION system privilege, 7-522
- ALTER ANY INDEX system privilege, 7-522
- ALTER ANY INDEXTYPE system privilege, 7-522
- ALTER ANY MATERIALIZED VIEW system privilege, 7-523
- ALTER ANY OUTLINE system privilege, 7-524
- ALTER ANY PROCEDURE system privilege, 7-524
- ALTER ANY ROLE system privilege, 7-525
- ALTER ANY SEQUENCE system privilege, 7-525
- ALTER ANY SNAPSHOT system privilege, 7-526
- ALTER ANY TABLE system privilege, 7-526
- ALTER ANY TRIGGER system privilege, 7-527
- ALTER ANY TYPE system privilege, 7-528
- ALTER CLUSTER statement, 7-2
- ALTER DATABASE
  - statement, 7-6
  - system privilege, 7-521
- ALTER DIMENSION statement, 7-25
- ALTER FUNCTION statement, 7-28
- ALTER INDEX statement, 7-30
- ALTER JAVA CLASS statement, 7-45
- ALTER JAVA SOURCE statement, 7-45
- ALTER MATERIALIZED VIEW LOG
  - statement, 7-58
- ALTER MATERIALIZED VIEW statement, 7-47
- ALTER object privilege, 7-535
- ALTER OUTLINE statement, 7-63
- ALTER PACKAGE statement, 7-64
- ALTER PROCEDURE statement, 7-67
- ALTER PROFILE
  - statement, 7-69
  - system privilege, 7-525
- ALTER RESOURCE COST
  - statement, 7-73
  - system privilege, 7-525
- ALTER ROLE statement, 7-76
- ALTER ROLLBACK SEGMENT
  - statement, 7-78
  - system privilege, 7-525
- ALTER SEQUENCE statement, 7-81
- ALTER SESSION
  - statement, 7-83
  - system privilege, 7-525
- ALTER SNAPSHOT LOG. *See* ALTER MATERIALIZED VIEW LOG
- ALTER SNAPSHOT. *See* ALTER MATERIALIZED VIEW
- ALTER statement
  - triggers on, 7-430
- ALTER SYSTEM
  - statement, 7-102
  - system privilege, 7-521
- ALTER TABLE statement, 7-123
- ALTER TABLESPACE
  - statement, 7-176
  - system privilege, 7-527



- ALTER TRIGGER statement, 7-183
- ALTER TYPE statement, 7-186
- ALTER USER
  - statement, 7-193
  - system privilege, 7-528
- ALTER VIEW statement, 7-198
- AM (A.M.) date format element, 2-46, 2-48
- American National Standards Institute. *See* ANSI
- analytic functions, 4-7
  - CUME\_DIST, 4-28
  - FIRST\_VALUE, 4-34
  - LAG, 4-40
  - LAST\_VALUE, 4-42
  - LEAD, 4-44
  - NTILE, 4-61
  - PERCENT\_RANK, 4-66
  - RANK, 4-67
  - RATIO\_TO\_REPORT, 4-68
  - ROW\_NUMBER, 4-80
- ANALYZE ANY system privilege, 7-529
- ANALYZE CLUSTER statement, 7-200
- ANALYZE INDEX statement, 7-200
- ANALYZE TABLE statement, 7-200
- ANCILLARY TO clause
  - of CREATE OPERATOR, 7-341
- AND operator, 3-11, 3-12
- AND\_EQUAL hint, 2-65
- ANSI, B - 1
  - datatypes, 2-28
    - conversion to Oracle datatypes, 2-28
  - standards, xv, 1-2
  - supported datatypes, 2-12
- ANY operator, 3-6
- APPEND hint, 2-67
- application servers
  - allowing to connect as a user, 7-195
- applications
  - allowing to connect as a user, 7-195
  - securing, 7-261
  - validating, 7-261
- AQ\_ADMINISTRATOR\_ROLE role, 7-530
- AQ\_TM\_PROCESSES parameter
  - of ALTER SYSTEM, 7-108
- AQ\_USER\_ROLE role, 7-530
- ARCHIVE LOG clause
  - of ALTER SYSTEM, 7-103
- archived redo logs
  - location of, 7-13
  - storage locations, 7-89, 7-112
- ARCHIVELOG clause
  - of ALTER DATABASE, 7-9, 7-18
  - of CREATE CONTROLFILE, 7-266
- ARCHVIELOG clause
  - OF CREATE DATABASE, 7-270
- arguments of operators, 3-1
- arithmetic operators, 3-3
- AS 'filespec' clause
  - of CREATE LIBRARY, 7-317
- AS clause
  - of CREATE JAVA, 7-314
- AS EXTERNAL clause
  - of CREATE FUNCTION, 7-289, 7-357
  - of CREATE TYPE BODY, 7-450
- AS OBJECT clause
  - of CREATE TYPE, 7-440
- AS subquery
  - of CREATE MATERIALIZED VIEW / SNAPSHOT, 7-319, 7-328
  - of CREATE TABLE, 7-409
  - of CREATE VIEW, 7-459
- AS TABLE clause
  - of CREATE TYPE, 7-444
- AS VARRAY clause
  - of CREATE TYPE, 7-444
- ASC clause
  - of CREATE INDEX, 7-299
- ascending indexes, 7-299
- ASCII
  - character set, 2-34
- ASCII function, 4-14
- ASIN function, 4-15
- ASSOCIATE STATISTICS statement, 7-210
- ATAN function, 4-15
- ATAN2 function, 4-16
- ATTRIBUTE clause
  - of ALTER DIMENSION, 7-26
  - of CREATE DIMENSION, 7-280
- attributes
  - adding to a dimension, 7-27
  - dropping from a dimension, 7-27

- maximum number of in object type, 7-392
  - of dimensions, defining, 7-280
- AUDIT (Schema Objects) statement, 7-221
- AUDIT ANY system privilege, 7-529
- AUDIT sql\_statements statement, 7-213
- AUDIT SYSTEM system privilege, 7-521
- auditing
  - options
    - for database objects, 7-215
    - for SQL statements, 7-218
  - schema objects
    - stopping, 7-552
  - SQL statements, 7-215
  - SQL statements, stopping, 7-550
- AUTHENTICATED BY clause
  - of CREATE DATABASE LINK, 7-275
- AUTHID CURRENT\_USER clause
  - of ALTER JAVA, 7-46
  - of CREATE FUNCTION, 7-288
  - of CREATE JAVA, 7-313
  - of CREATE PACKAGE, 7-345
  - of CREATE PROCEDURE, 7-357
  - of CREATE TYPE, 7-188, 7-441
- AUTHID DEFINER clause
  - of ALTER JAVA, 7-46
  - of CREATE FUNCTION, 7-288
  - of CREATE JAVA, 7-313
  - of CREATE PACKAGE, 7-345
  - of CREATE PROCEDURE, 7-357
  - of CREATE TYPE, 7-188, 7-441
- AUTOEXTEND clause
  - for datafiles, 7-17
  - of ALTER DATABASE, 7-10
  - of ALTER TABLESPACE, 7-177, 7-179
  - of CREATE DATABASE, 7-268
  - of CREATE TABLESPACE, 7-419, 7-421
  - of CREATE TEMPORARY
    - TABLESPACE, 7-424, 7-425
- AVG function, 4-16
- AY date format element, 2-46

**B**

---

- BACKGROUND\_DUMP\_DEST parameter
  - of ALTER SYSTEM, 7-108
- BACKUP ANY TABLE system privilege, 7-526
- BACKUP CONTROLFILE clause
  - of ALTER DATABASE, 7-9, 7-20
- BACKUP\_TAPE\_IO\_SLAVES parameter
  - of ALTER SYSTEM, 7-108
- BC (B.C.) date format element, 2-46, 2-48
- BECOME USER system privilege, 7-528
- BEFORE clause
  - of CREATE TRIGGER, 7-428
- BEFORE triggers, 7-428
- BEGIN BACKUP clause
  - of ALTER TABLESPACE, 7-180
- BFILE
  - datatype, 2-25
  - locators, 2-25
- BFILENAME function, 4-17
- binary large objects. *See* BLOBs
- binary operators, 3-1
- BINDING clause
  - of CREATE OPERATOR, 7-339, 7-340
- BITMAP clause
  - of CREATE INDEX, 7-297
- bitmap indexes, 7-297
- blank padding
  - specifying in format models, 2-51
  - suppressing, 2-51
- BLOB datatype, 2-25
  - transactional support of, 2-25
- BODY clause
  - of ALTER PACKAGE, 7-65
- BUFFER\_POOL parameter
  - of STORAGE clause, 7-609
- BUILD DEFERRED clause
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-323
- BUILD IMMEDIATE clause
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-323
- BY ACCESS clause
  - of AUDIT schema\_objects, 7-222
  - of AUDIT sql\_statements, 7-215
- BY proxy clause
  - of AUDIT (SQL statements), 7-214
  - of NOAUDIT sql\_statements, 7-551
- BY SESSION clause

- of AUDIT schema\_objects, 7-222
- of AUDIT sql\_statements, 7-214
- BY user clause
  - of AUDIT sql\_statements, 7-214
  - of NOAUDIT sql\_statements, 7-551

## C

---

- C clause
  - of CREATE TYPE, 7-442
  - of CREATE TYPE BODY, 7-450
- C method
  - mapping to an object type, 7-442
- C number format element, 2-41
- CACHE clause
  - of ALTER MATERIALIZED VIEW, 7-51
  - of ALTER MATERIALIZED VIEW LOG, 7-61
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE, 7-81
  - of ALTER TABLE, 7-151
  - of CREATE CLUSTER, 7-258
  - of CREATE MATERIALIZED VIEW LOG/SNAPSHOT LOG, 7-336
  - of CREATE MATERIALIZED VIEW/SNAPSHOT, 7-323
  - of CREATE SEQUENCE, 7-374
  - of CREATE TABLE, 7-408
- CACHE hint, 2-67
- CALL clause
  - of CREATE TRIGGER, 7-433
- CALL procedure statement
  - of CREATE TRIGGER, 7-433
- call spec
  - in procedures, 7-354
  - of CREATE FUNCTION, 7-289
  - of CREATE PROCEDURE, 7-357
  - of CREATE TYPE, 7-442
  - of CREATE TYPE BODY, 7-450
- call specifications. *See* call spec
- CALL statement, 7-226
- Cartesian products, 5-23
- CASCADE clause
  - of CREATE TABLE, 7-408
  - of DROP PROFILE, 7-494
  - of DROP USER, 7-510
- CASCADE CONSTRAINTS clause
  - of DROP CLUSTER, 7-473
  - of DROP TABLE, 7-502
  - of DROP TABLESPACE, 7-504
  - of REVOKE schema\_object\_privileges, 7-561
- CASE expressions, 5-13
- CAST expressions, 5-7
- CC date format element, 2-46
- CEIL function, 4-18
- century
  - specifying, 2-47
- chained rows
  - listing, 7-207
- CHANGE CATEGORY clause
  - of ALTER OUTLINE, 7-63
- changes
  - making permanent, 7-230
- changing default storage parameters, 7-180
- CHAR datatype, 2-15
  - ANSI, 2-28
  - converting to VARCHAR2, 2-40
- CHAR VARYING datatype, ANSI, 2-28
- CHARACTER datatype
  - ANSI, 2-28
  - DB2, 2-29
  - SQL/DS, 2-29
- character functions, 4-4, 4-5
- character large objects. *See* CLOB datatype
- character literal. *See* text
- CHARACTER SET clause
  - of CREATE CONTROLFILE, 7-266
  - OF CREATE DATABASE, 7-271
- CHARACTER SET parameter
  - of ALTER DATABASE, 7-16
- character sets
  - common, 2-33
  - multibyte characters, 2-72
  - specifying for database, 7-271
- character strings
  - comparison rules, 2-32
  - exact matching of, 2-51
  - fixed-length, 2-15
  - national character set, 2-15
  - variable length, 2-16
  - variable-length, 2-19

- zero-length, 2-15
- CHARACTER VARYING datatype
  - ANSI, 2-28
- characters
  - single, comparison rules, 2-33
- CHARTOROWID function, 2-37, 4-19
- CHECK clause
  - of constraint\_clause, 7-241
  - of CREATE TABLE, 7-393
- check constraints, 7-241
- CHECK DATAFILES clause
  - of ALTER SYSTEM, 7-105
- checkpoint
  - forcing, 7-105
- CHECKPOINT clause
  - of ALTER SYSTEM, 7-105
- CHOOSE hint, 2-64
- CHR function, 4-19
- CHUNK clause
  - of ALTER TABLE, 7-141
  - of CREATE TABLE, 7-399
- CLEAR LOGFILE clause
  - of ALTER DATABASE, 7-9, 7-19
- CLOB datatype, 2-25
  - transactional support of, 2-25
- clone database
  - mounting, 7-11
- CLOSE DATABASE LINK clause
  - of ALTER SESSION, 7-84
- CLUSTER clause
  - of CREATE INDEX, 7-297
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-322
  - of CREATE TABLE, 7-401
  - of TRUNCATE, 7-613
- CLUSTER hint, 2-65
- clusters
  - allocating extents for, 7-3
  - assigning tables to, 7-401
  - caching retrieved blocks, 7-258
  - cluster indexes on, 7-297
  - collecting statistics on, 7-204
  - creating, 7-254, 7-255
  - data blocks allocated to, 7-256
  - deallocating unused extents, 7-3
  - degree of parallelism
    - changing, 7-3
    - when creating, 7-258
  - dropping tables of, 7-473
  - granting
    - system privileges on, 7-521
  - hash, 7-257
    - single-table, 7-257
  - indexed, 7-257
  - key
    - changing column names, 7-3
    - changing the number of columns, 7-3
    - migrated and chained rows in, 7-207
    - modifying, 7-2
    - physical attributes
      - changing, 7-2
      - specifying, 7-256
    - removing from the database, 7-472
    - space allocated for cluster key values, 7-256
    - SQL examples, 7-473
    - storage attributes
      - changing, 7-2
    - storage characteristics, 7-605
      - specifying, 7-256
    - tablespace
      - changing, 7-3
      - tablespace in which created, 7-257
      - validating structure of, 7-206
  - COALESCE clause
    - for partitions, 7-160
    - for subpartitions, 7-156
    - of ALTER INDEX, 7-40
    - of ALTER TABLESPACE, 7-181
  - COALESCE SUBPARTITION clause
    - of ALTER TABLE, 7-156
  - code examples
    - description of, xxiii
  - collections
    - inserting rows into, 7-542
    - modifying, 7-145
    - nested tables, 2-31
    - treating as a table, 7-467, 7-542, 7-617
    - unnesting, 7-576
      - examples, 7-593
    - varrays, 2-31

- column constraints, 7-233, 7-237
  - of ALTER TABLE, 7-141
  - of CREATE TABLE, 7-393
- column REF constraints, 7-234, 7-242
  - of ALTER TABLE, 7-140
  - of CREATE TABLE, 7-393
- columns
  - adding, 7-140
  - aliases for, 5-19
  - associating statistics with, 7-211
  - basing an index on, 7-298
  - collecting statistics on, 7-204
  - creating comments about, 7-228
  - defining, 7-390
  - LOB, storage characteristics of, 7-141
  - maximum number of, 7-392
  - modifying existing, 7-143
  - parent-child relationships between, 7-278
  - prohibiting nulls in, 7-239
  - qualifying names of, 5-19
  - REF
    - describing, 7-242
    - restricting values for, 7-237
    - specifying as foreign key, 7-240
    - specifying as primary key, 7-239
    - specifying constraints on, 7-393
    - specifying default values for, 7-392
    - unique values in, 7-238
- COLUMNS clause
  - of ASSOCIATE STATISTICS, 7-210, 7-211
- COMMENT ANY TABLE system privilege, 7-529
- COMMENT clause
  - of COMMIT, 7-231
- COMMENT statement, 7-228
- comments, 2-62
  - adding to objects, 7-228
  - associating with a transaction, 7-231
  - dropping from objects, 7-228
  - how to specify, 2-62
  - in SQL statements, 2-62
  - on schema objects, 2-63
  - removing from the data dictionary, 7-228
  - viewing, 7-228
- commit
  - automatic, 7-230
  - COMMIT IN PROCEDURE clause
    - of ALTER SESSION, 7-84
  - COMMIT statement, 7-230
  - comparison functions
    - MAP, 7-443, 7-449
    - ORDER, 7-443, 7-449
  - comparison operators, 3-5
  - comparison semantics
    - blank-padded, 2-33
    - nonpadded, 2-33
    - of character strings, 2-32
  - COMPILE clause
    - of ALTER DIMENSION, 7-27
    - of ALTER FUNCTION, 7-28
    - of ALTER JAVA SOURCE, 7-46
    - of ALTER MATERIALIZED VIEW, 7-54
    - of ALTER PACKAGE, 7-65
    - of ALTER PROCEDURE, 7-68
    - of ALTER TRIGGER, 7-184
    - of ALTER TYPE, 7-187
    - of ALTER VIEW, 7-199
    - of CREATE JAVA, 7-312
  - compiler directives, 7-442
  - composite foreign keys, 7-239
  - composite partitioning clause
    - of CREATE TABLE, 7-386, 7-403
  - composite primary keys, 7-239
  - composite unique constraints, 7-238
  - COMPOSITE\_LIMIT parameter
    - of ALTER PROFILE, 7-69
    - of CREATE PROFILE, 7-362
  - compound conditions, 5-18
  - compound expressions, 5-4
  - COMPRESS clause
    - of ALTER INDEX, 7-32
    - of ALTER TABLE, 7-144
    - of CREATE INDEX, 7-300
    - of CREATE TABLE, 7-397
  - COMPUTE STATISTICS clause
    - of ANALYZE, 7-204
    - of CREATE INDEX, 7-301
  - CONCAT function, 4-20
  - concatenation operator, 3-3
  - conditions
    - compound, 5-18

- EXISTS, 5-18
  - group comparison, 5-16
  - in SQL syntax, 5-14
- LIKE, 5-18
  - membership, 5-17
- NULL, 5-18
  - range, 5-18
  - simple comparison, 5-16
- CONNECT BY clause
  - of SELECT, 5-21, 7-577
- CONNECT clause
  - of SELECT and subqueries, 7-571
- CONNECT role, 7-530
- CONNECT TO clause
  - of CREATE DATABASE LINK, 7-274
- CONNECT\_TIME parameter
  - of ALTER PROFILE, 7-69
  - of ALTER RESOURCE COST, 7-74
  - of CREATE PROFILE, 7-361
- CONSIDER FRESH clause
  - of ALTER MATERIALIZED VIEW, 7-55
- constant values. *See* literals
- CONSTRAINT clause
  - of constraint\_clause, 7-237
- constraint clause, 7-233
- CONSTRAINT(S) parameter
  - of ALTER SESSION, 7-86
- constraints
  - adding, 7-140
  - check, 7-241
  - checking at end of transaction, 7-243
  - checking at start of transaction, 7-243
  - checking at the end of each DML
    - statement, 7-243
  - column REF, 7-242
  - composite unique, 7-238
  - deferrable, 7-243, 7-598
    - enforcing, 7-86
  - defining, 7-237, 7-390
    - on a column, 7-393
    - on a table, 7-393
  - disabling, 7-165, 7-245, 7-405
    - cascading, 7-408
  - dropping, 7-146, 7-504
  - enabling, 7-165, 7-244, 7-405, 7-407
    - foreign key, 7-240
    - modifying existing, 7-144
    - not null, 7-239
    - on columns, 7-237
    - primary key, 7-239
      - attributes of index, 7-244
      - enabling, 7-407
    - referential integrity, 7-239, 7-240
    - restrictions, 7-237
    - scope, 7-242
    - setting state for a transaction, 7-598
    - storing rows in violation, 7-163, 7-245, 7-407
  - table REF, 7-242
  - unique, 7-238
    - attributes of index, 7-244
    - composite, 7-238
    - enabling, 7-407
    - validating, 7-244, 7-245
- constructor methods
  - and object types, 7-439
- context namespaces
  - removing from the database, 7-474
- contexts
  - creating namespaces for, 7-261
  - granting
    - system privileges on, 7-521
  - namespace
    - associating with package, 7-261
- control files
  - allow reuse of, 7-264
  - allowing reuse of, 7-269
  - backing up, 7-20
  - re-creating, 7-263
- CONTROL\_FILE\_RECORD\_KEEP\_TIME parameter
  - of ALTER SYSTEM, 7-108
- controlfile clauses
  - of ALTER DATABASE, 7-9
- CONTROLFILE REUSE clause
  - of CREATE DATABASE, 7-269
- conversion
  - rules, string to date, 2-53
- conversion functions, 4-5
  - table of, 2-37
- CONVERT clause
  - of ALTER DATABASE, 7-11

CONVERT function, 4-21  
 CORE\_DUMP\_DEST parameter  
     of ALTER SYSTEM, 7-108  
 CORR function, 4-22  
 correlated subqueries, 5-25  
 correlation names  
     for base tables of indexes, 7-298  
     in DELETE, 7-467  
     in SELECT, 7-576  
 COS function, 4-23  
 COSH function, 4-24  
 COUNT function, 4-24  
 CPU\_PER\_CALL parameter  
     of ALTER PROFILE, 7-69  
     of CREATE PROFILE, 7-361  
 CPU\_PER\_SESSION parameter  
     of ALTER PROFILE, 7-69  
     of ALTER RESOURCE COST, 7-74  
     of CREATE PROFILE, 7-361  
 CREATE ANY CLUSTER system privilege, 7-521  
 CREATE ANY CONTEXT system privilege, 7-521  
 CREATE ANY DIMENSION system  
     privilege, 7-522  
 CREATE ANY DIRECTORY system  
     privilege, 7-522  
 CREATE ANY INDEX system privilege, 7-522  
 CREATE ANY INDEXTYPE system  
     privilege, 7-522  
 CREATE ANY LIBRARY system privilege, 7-523  
 CREATE ANY MATERIALIZED VIEW system  
     privilege, 7-523  
 CREATE ANY OPERATOR system  
     privilege, 7-524  
 CREATE ANY OUTLINE system privilege, 7-524  
 CREATE ANY PROCEDURE system  
     privilege, 7-524  
 CREATE ANY SEQUENCE system  
     privilege, 7-525  
 CREATE ANY SNAPSHOT system  
     privilege, 7-526  
 CREATE ANY SYNONYM system privilege, 7-526  
 CREATE ANY TABLE system privilege, 7-526  
 CREATE ANY TRIGGER system privilege, 7-527  
 CREATE ANY TYPE system privilege, 7-528  
 CREATE ANY VIEW system privilege, 7-529  
 CREATE CLUSTER  
     statement, 7-254  
     system privilege, 7-521  
 CREATE CONTEXT statement, 7-261  
 CREATE CONTROLFILE statement, 7-263  
 CREATE DATABASE LINK  
     statement, 7-273  
     system privilege, 7-522  
 CREATE DATABASE statement, 7-267  
 CREATE DATAFILE clause  
     of ALTER DATABASE, 7-8, 7-16  
 CREATE DIMENSION  
     statement, 7-277  
     system privilege, 7-522  
 CREATE DIRECTORY statement, 7-282  
 CREATE FUNCTION statement, 7-284  
 CREATE INDEX  
     statement, 7-291  
 CREATE INDEXTYPE  
     statement, 7-309  
     system privilege, 7-522  
 CREATE JAVA statement, 7-311  
 CREATE LIBRARY  
     statement, 7-316  
     system privilege, 7-523  
 CREATE MATERIALIZED VIEW / SNAPSHOT  
     statement, 7-318  
 CREATE MATERIALIZED VIEW LOG /  
     SNAPSHOT LOG statement, 7-333  
 CREATE MATERIALIZED VIEW / SNAPSHOT  
     system privilege, 7-523  
 CREATE OPERATOR  
     statement, 7-339  
     system privilege, 7-524  
 CREATE OUTLINE statement, 7-342  
 CREATE PACKAGE BODY statement, 7-348  
 CREATE PACKAGE statement, 7-344  
 CREATE PROCEDURE  
     statement, 7-353  
     system privilege, 7-524  
 CREATE PROFILE  
     statement, 7-359  
     system privilege, 7-524  
 CREATE PUBLIC DATABASE LINK system  
     privilege, 7-522

- CREATE PUBLIC SYNONYM system privilege, 7-526
- CREATE ROLE
  - statement, 7-365
  - system privilege, 7-525
- CREATE ROLLBACK SEGMENT
  - statement, 7-367
  - system privilege, 7-525
- CREATE SCHEMA statement, 7-369
- CREATE SEQUENCE
  - statement, 7-371
  - system privilege, 7-525
- CREATE SESSION system privilege, 7-525
- CREATE SNAPSHOT system privilege, 7-526
- CREATE STANDBY CONTROLFILE clause of ALTER DATABASE, 7-9, 7-20
- CREATE statement
  - triggers on, 7-430
- CREATE SYNONYM
  - statement, 7-377
  - system privilege, 7-526
- CREATE TABLE statement, 7-381
- CREATE TABLESPACE
  - statement, 7-419
  - system privilege, 7-527
- CREATE TEMPORARY TABLESPACE
  - statement, 7-424
- CREATE TRIGGER
  - statement, 7-426
  - system privilege, 7-527
- CREATE TYPE
  - statement, 7-437
  - system privilege, 7-528
- CREATE TYPE BODY statement, 7-447
- CREATE USER
  - statement, 7-451
  - system privilege, 7-528
- CREATE VIEW
  - statement, 7-456
  - system privilege, 7-529
- CREATE\_STORED\_OUTLINES parameter
  - of ALTER SESSION, 7-86
  - of ALTER SYSTEM, 7-109
- cross-tabulation values
  - deriving, 7-578

- CUBE operations
  - of queries and subqueries, 7-578
- CUME\_DIST function, 4-28
- currency symbol
  - ISO, 2-41
  - local, 2-42
  - union, 2-43
- CURRENT\_SCHEMA parameter of ALTER SESSION, 7-87
- CURRENT\_USER
  - and database links, 7-274
- CURRVAL pseudocolumn, 2-56, 7-372
- CURSOR expressions, 5-10
- CURSOR\_SHARING parameter of ALTER SESSION, 7-87, 7-109
- cursors
  - number cached per session, 7-94
- CYCLE clause
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE, 7-81
  - of CREATE SEQUENCE, 7-373

## D

---

- D date format element, 2-46
- D number format element, 2-41
- data
  - integrity checking on input, 2-18
  - retrieving, 5-19
  - undo
    - storing, 7-367
- data conversion, 2-36
  - implicit versus explicit, 2-38
  - when performed implicitly, 2-36
  - when specified explicitly, 2-37
- data definition language
  - events and triggers, 7-430
  - statements, 6-1
    - and implicit commit, 6-2
    - causing recompilation, 6-2
    - PL/SQL support of, 6-2
- data dictionary
  - adding comments to, 7-228
- data manipulation language
  - operations



- and triggers, 7-429
  - during index creation, 7-301
  - during index rebuild, 7-144
- statements, 6-4
  - PL/SQL support of, 6-4
- data object number
  - in extended rowids, 2-27
- database
  - accounts
    - creating, 7-451
  - allowing generation of redo logs, 7-12
  - allowing reuse of control files, 7-269
  - allowing unlimited resources to users, 7-361
  - cancel-based recovery, 7-13
    - terminating, 7-14
  - change-based recovery, 7-14
  - changing characteristics of, 7-263
  - changing global name, 7-15
  - changing the name of, 7-263, 7-264
  - character set
    - specifying, 7-271
  - connect strings, 2-80
  - converting from Oracle7 data dictionary, 7-11
  - creating, 7-268
  - designing media recovery, 7-12
  - enabling automatic extension of, 7-271
  - erasing all data from, 7-268
  - granting system privileges on, 7-521
  - limiting resources for users, 7-360
  - managed recovery of, 7-8
  - modifying, 7-10
  - mounting, 7-11, 7-268
  - naming, 7-11
  - opening, 7-11, 7-268
    - after media recovery, 7-12
  - recovering, 7-13
    - with backup control file, 7-14
  - re-creating control file for, 7-263
  - redo log files
    - specifying, 7-265
  - remote
    - accessing, 5-26
    - authenticating users to, 7-275
    - connecting to, 7-274
    - inserting into, 7-542
    - service name of, 7-275
    - table locks on, 7-548
  - resetting
    - current log sequence, 7-12
    - to an earlier version, 7-15
  - restricting users to read-only transactions, 7-12
  - resuming activity, 7-108
  - specifying datafiles for, 7-265
  - suspending activity, 7-107
  - time-based recovery, 7-13
- database events
  - and triggers, 7-431
- database links, 5-26
  - closing, 7-84
  - creating, 2-79, 7-273
  - creating synonyms with, 7-379
  - current user, 7-274
  - granting system privileges on, 7-521
  - naming, 2-79
  - public, 7-274
    - dropping, 7-475
  - referring to, 2-81
  - removing from the database, 7-475
  - shared, 7-274
  - syntax of, 2-80
  - username and password, 2-80
- database objects
  - dropping, 7-510
  - nonschema, 2-69
  - schema, 2-68
- database triggers. *See* triggers
- DATAFILE clause
  - of ALTER DATABASE, 7-8, 7-16
  - of CREATE CONTROLFILE, 7-265
  - of CREATE DATABASE, 7-271
- DATAFILE clauses
  - of ALTER DATABASE, 7-8
- DATAFILE END BACKUP clause
  - of ALTER DATABASE, 7-17
- DATAFILE OFFLINE clause
  - of ALTER DATABASE, 7-17
- DATAFILE ONLINE clause
  - of ALTER DATABASE, 7-17
- DATAFILE RESIZE clause
  - of ALTER DATABASE, 7-17

- datafiles
  - bringing online, 7-17
  - creating new, 7-16
  - designing media recovery, 7-12
  - disabling automatic extension, 7-17
  - enabling automatic extension, 7-17, 7-421
  - modifying, 7-16
  - recovering, 7-14
  - re-creating lost, 7-16
  - renaming, 7-15
  - resizing, 7-17
  - reusing, 7-517
  - size of, 7-517
  - specifying, 7-516
    - for a tablespace, 7-421
  - taking offline, 7-17
- datatypes, 2-9
  - ANSI-supported, 2-12
  - associating statistics with, 7-212
  - BFILE, 2-14, 2-25
  - BLOB, 2-14, 2-25
  - built-in, 2-13
    - syntax, 2-11
  - CHAR, 2-13, 2-15
  - character, 2-14
  - CLOB, 2-14, 2-25
  - comparison rules, 2-32
  - conversion
    - table of, 2-37
  - DATE, 2-13, 2-21
  - external, 2-10
  - LONG, 2-13, 2-19
  - LONG RAW, 2-13, 2-22
  - NCHAR, 2-14, 2-15
  - NCLOB, 2-14, 2-26
  - NUMBER, 2-17
  - NUMER, 2-13
  - NVARCHAR2, 2-13, 2-16
  - RAW, 2-13, 2-22
  - ROWID, 2-13, 2-26
  - UROWID, 2-13, 2-27
  - VARCHAR, 2-17
  - VARCHAR2, 2-13, 2-16
- DATE datatype, 2-21
  - converting from character or numeric
    - value, 2-21
- date format elements, 2-45
  - and NLS, 2-48
  - capitalization, 2-45
  - ISO standard, 2-49
  - RR, 2-49
  - suffixes, 2-50
- date format models, 2-45
  - punctuation in, 2-45
  - text in, 2-45
- date functions, 4-5
- dates
  - arithmetic using, 2-21
  - comparison rules, 2-32
  - Julian, 2-22
- DAY date format element, 2-48
- DB\_BLOCK\_CHECKING parameter
  - of ALTER SESSION, 7-87
  - of ALTER SYSTEM, 7-109
- DB\_BLOCK\_CHECKSUM parameter
  - of ALTER SYSTEM, 7-109
- DB\_BLOCK\_MAX\_DIRTY\_TARGET parameter
  - of ALTER SYSTEM, 7-109
- DB\_FILE\_DIRECT\_IO\_COUNT parameter
  - of ALTER SYSTEM, 7-110
- DB\_FILE\_MULTIBLOCK\_READ\_COUNT
  - parameter
    - of ALTER SESSION, 7-87
    - of ALTER SYSTEM, 7-110
- DB2 datatypes, 2-28
  - conversion to Oracle datatypes, 2-29
  - restrictions on, 2-30
- DBA role, 7-530
- DBA\_2PC\_PENDING view, 7-84
- DBA\_COL\_COMMENTS view, 7-228
- DBA\_ROLLBACK\_SEGS view, 7-496
- DBA\_TAB\_COMMENTS view, 7-228
- DBMS\_OUTPUT package, 7-184
- DBMS\_ROWID package
  - and extended rowids, 2-27
- DBMSSTD.SQL script, 7-285, 7-344, 7-348, 7-355
  - and triggers, 7-427
- DD date format element, 2-46
- DDD date format element, 2-46
- DDL. *See* data definition language

- DDL statements
  - requiring exclusive access, 6-2
- DEALLOCATE UNUSED clause
  - of ALTER CLUSTER, 7-3, 7-5
  - of ALTER INDEX, 7-30
  - of ALTER TABLE, 7-150
- DEBUG clause
  - of ALTER FUNCTION, 7-28
  - of ALTER PACKAGE, 7-65
  - of ALTER PROCEDURE, 7-68
  - of ALTER TRIGGER, 7-184
  - of ALTER TYPE, 7-188
- decimal characters, 2-4
  - specifying, 2-42
- DECIMAL datatype
  - ANSI, 2-28
  - DB2, 2-29
  - SQL/DS, 2-29
- DECODE expressions, 5-12
- DEFAULT clause
  - of CREATE TABLE, 7-392
- DEFAULT COST clause
  - of ASSOCIATE STATISTICS, 7-211, 7-212
- DEFAULT profile
  - assigning to users, 7-494
- DEFAULT ROLE clause
  - of ALTER USER, 7-195
- DEFAULT SELECTIVITY clause
  - of ASSOCIATE STATISTICS, 7-211, 7-212
- DEFAULT storage clause
  - of ALTER TABLESPACE, 7-180
  - of CREATE TABLESPACE, 7-422
- DEFAULT TABLESPACE clause
  - of ALTER USER. *See* CREATE USER
  - of CREATE USER, 7-453
- DEFERRABLE clause
  - of constraint\_clause, 7-243
- deferrable constraints, 7-598
- DEFERRED clause
  - of SET CONSTRAINTS, 7-598
- DELETE
  - object privilege, 7-535
  - statement, 7-464
- DELETE ANY TABLE system privilege, 7-526
- DELETE statement
  - triggers on, 7-429
- DELETE STATISTICS clause
  - of ANALYZE, 7-206
- DELETE\_CATALOG\_ROLE role, 7-530
- DENSE\_RANK function, 4-29
- DEREF function, 4-30
- DESC clause
  - of CREATE INDEX, 7-299
- descending indexes, 7-299
- DETERMINISTIC clause
  - of CREATE FUNCTION, 7-288
- dimensions
  - attributes
    - adding, 7-27
    - changing, 7-26
    - defining, 7-280
    - dropping, 7-27
  - changing hierarchical relationships, 7-26
  - compiling invalidated, 7-27
  - creating, 7-278
  - examples, 7-280
  - granting
    - system privileges on, 7-522
  - hierarchies
    - adding, 7-27
    - defining, 7-279
    - dropping, 7-27
  - levels
    - adding, 7-27
    - defining, 7-279
    - dropping, 7-27
  - removing from the database, 7-476
- directories. *See* directory objects
- directory objects
  - as aliases for OS directories, 7-282
  - auditing, 7-222
  - creating, 7-282
  - granting system privileges on, 7-522
  - redefining, 7-283
  - removing from the database, 7-477
- DISABLE [constraint] clause
  - of CREATE TABLE, 7-406
- DISABLE ALL TRIGGERS clause
  - of ALTER TABLE, 7-165
- DISABLE clause

- of ALTER INDEX, 7-40
- of ALTER TRIGGER, 7-184
- of constraint\_clause, 7-245
- of CREATE TABLE, 7-405
- DISABLE DISTRIBUTED RECOVERY clause
  - of ALTER SYSTEM, 7-106
- DISABLE NOVALIDATE constraint state, 7-407
- DISABLE PARALLEL DML clause
  - of ALTER SESSION, 7-84
- DISABLE QUERY REWRITE clause
  - of ALTER MATERIALIZED VIEW, 7-54
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-327
- DISABLE RESTRICTED SESSION clause
  - of ALTER SYSTEM, 7-106
- DISABLE ROW MOVEMENT clause
  - of ALTER TABLE, 7-164
  - of CREATE TABLE, 7-383, 7-404
- DISABLE STORAGE IN ROW clause
  - of ALTER TABLE, 7-141
  - of CREATE TABLE, 7-398
- DISABLE TABLE LOCK clause
  - of ALTER TABLE, 7-165
- DISABLE THREAD clause
  - of ALTER DATABASE, 7-16
- DISABLE VALIDATE constraint state, 7-406
- DISASSOCIATE STATISTICS statement, 7-470
- DISCONNECT SESSION clause
  - of ALTER SYSTEM, 7-106
- dispatcher processes
  - creating additional, 7-114
  - terminating, 7-114
- DISTINCT clause
  - of SELECT, 7-573
- distinct queries, 7-573
- distributed queries, 5-26
  - restrictions on, 5-26
- DML. *See* data manipulation language
- domain indexes, 7-295, 7-304, 7-309
  - associating statistics with, 7-212
  - determining user-defined CPU and I/O  
costs, 7-512
  - invoking drop routines for, 7-501
  - removing from the database, 7-480
  - specifying alter string for, 7-40

- DOUBLE PRECISION datatype
  - ANSI, 2-28
- DRIVING\_SITE hint, 2-66
- DROP ANY CLUSTER system privilege, 7-521
- DROP ANY CONTEXT system privilege, 7-521
- DROP ANY DIMENSION system privilege, 7-522
- DROP ANY DIRECTORY system privilege, 7-522
- DROP ANY INDEX system privilege, 7-522
- DROP ANY INDEXTYPE system privilege, 7-522
- DROP ANY LIBRARY system privilege, 7-523
- DROP ANY MATERIALIZED VIEW system  
privilege, 7-523
- DROP ANY OPERATOR system privilege, 7-524
- DROP ANY OUTLINE system privilege, 7-524
- DROP ANY PROCEDURE system privilege, 7-524
- DROP ANY ROLE system privilege, 7-525
- DROP ANY SEQUENCE system privilege, 7-525
- DROP ANY SNAPSHOT system privilege, 7-526
- DROP ANY SYNONYM system privilege, 7-526
- DROP ANY TABLE system privilege, 7-527
- DROP ANY TRIGGER system privilege, 7-527
- DROP ANY TYPE system privilege, 7-528
- DROP ANY VIEW system privilege, 7-529
- DROP clause
  - of ALTER DIMENSION, 7-27
- DROP CLUSTER statement, 7-472
- DROP COLUMN clause
  - of ALTER TABLE, 7-147
- DROP CONSTRAINT clause
  - of ALTER TABLE, 7-146
- DROP CONTEXT statement, 7-474
- DROP DATABASE LINK statement, 7-475
- DROP DIMENSION statement, 7-476
- DROP DIRECTORY statement, 7-477
- DROP FUNCTION statement, 7-478
- DROP INDEX statement, 7-480
- DROP INDEXTYPE statement, 7-482
- DROP JAVA statement, 7-483
- DROP LIBRARY
  - statement, 7-484
  - system privilege, 7-523
- DROP LOGFILE clause
  - of ALTER DATABASE, 7-9, 7-19
- DROP LOGFILE MEMBER clause
  - of ALTER DATABASE, 7-9, 7-19

DROP MATERIALIZED VIEW / SNAPSHOT  
 statement, 7-485

DROP MATERIALIZED VIEW LOG / SNAPSHOT  
 LOG statement, 7-487

DROP OPERATOR statement, 7-489

DROP OUTLINE statement, 7-490

DROP PACKAGE BODY statement, 7-491

DROP PACKAGE statement, 7-491

DROP PARTITION clause  
 of ALTER INDEX, 7-41  
 of ALTER TABLE, 7-160

DROP PRIMARY constraint clause  
 of ALTER TABLE, 7-146

DROP PROCEDURE statement, 7-493

DROP PROFILE  
 statement, 7-494  
 system privilege, 7-525

DROP PUBLIC DATABASE LINK system  
 privilege, 7-522

DROP PUBLIC SYNONYM system  
 privilege, 7-526

DROP ROLE statement, 7-495

DROP ROLLBACK SEGMENT  
 statement, 7-496  
 system privilege, 7-525

DROP SEQUENCE statement, 7-497

DROP statement  
 triggers on, 7-430

DROP STORAGE clause  
 of TRUNCATE, 7-613

DROP SYNONYM statement, 7-500

DROP TABLE statement, 7-501

DROP TABLESPACE  
 statement, 7-503  
 system privilege, 7-527

DROP TRIGGER statement, 7-505

DROP TYPE BODY statement, 7-508

DROP TYPE statement, 7-506

DROP UNIQUE constraint clause  
 of ALTER TABLE, 7-146

DROP USER  
 statement, 7-509  
 system privilege, 7-528

DROP VIEW statement, 7-511

DUAL dummy table, 2-73, 5-26

DUMP function, 4-31

DY date format element, 2-46, 2-48

## E

---

E date format element, 2-46

E number format element, 2-41

EBCDIC character set, 2-34

EE date format element, 2-46

embedded SQL, xv, 1-4, 6-5  
 precompiler support of, 6-5

EMPTY\_BLOB function, 4-33

EMPTY\_CLOB function, 4-33

ENABLE ALL TRIGGERS clause  
 of ALTER TABLE, 7-165

ENABLE clause  
 of ALTER INDEX, 7-40  
 of ALTER TRIGGER, 7-184  
 of constraint\_clause, 7-244  
 of CREATE TABLE, 7-405

ENABLE DISTRIBUTED RECOVERY clause  
 of ALTER SYSTEM, 7-106

ENABLE NOVALIDATE constraint state, 7-406

ENABLE PARALLEL DML clause  
 of ALTER SESSION, 7-84

ENABLE QUERY REWRITE clause  
 of ALTER MATERIALIZED VIEW, 7-54  
 of CREATE MATERIALIZED VIEW/  
 SNAPSHOT, 7-327

ENABLE RESTRICTED SESSION clause  
 of ALTER SYSTEM, 7-106

ENABLE ROW MOVEMENT clause  
 of ALTER TABLE, 7-164  
 of CREATE TABLE, 7-383, 7-404

ENABLE STORAGE IN ROW clause  
 of ALTER TABLE, 7-141  
 of CREATE TABLE, 7-398

ENABLE TABLE LOCK clause  
 of ALTER TABLE, 7-165

ENABLE THREAD clause  
 of ALTER DATABASE, 7-15

ENABLE VALIDATE constraint state, 7-406

ENABLE/DISABLE clause  
 of ALTER TABLE, 7-137, 7-165  
 of CREATE TABLE, 7-388

- END BACKUP clause
  - of ALTER TABLESPACE, 7-180
- equality test, 3-5
- equijoins, 5-22
  - defining for a dimension, 7-280
- equivalency tests, 3-6
- ESTIMATE STATISTICS clause
  - of ANALYZE, 7-204
- EXCEPTIONS INTO clause
  - of ALTER TABLE, 7-163
  - of constraint\_clause, 7-245
  - of CREATE TABLE, 7-407
- EXCHANGE PARTITION clause
  - of ALTER TABLE, 7-162
- EXCHANGE SUBPARTITION clause
  - of ALTER TABLE, 7-162
- EXCLUDING NEW VALUES clause
  - of ALTER MATERIALIZED VIEW LOG, 7-61
  - of CREATE MATERIALIZED VIEW LOG /  
SNAPSHOT LOG, 7-337
- EXCLUSIVE lock mode, 7-548
- EXECUTE ANY INDEXTYPE system
  - privilege, 7-522
- EXECUTE ANY OPERATOR system
  - privilege, 7-524
- EXECUTE ANY PROCEDURE system
  - privilege, 7-524
- EXECUTE ANY TYPE system privilege, 7-528
- EXECUTE object privilege, 7-535
- EXECUTE\_CATALOG\_ROLE role, 7-530
- execution plans
  - determining, 7-512
  - dropping outlines for, 7-490
  - saving, 7-342
- EXISTS
  - conditions, 5-18
  - operator, 3-6
- EXP function, 4-33
- EXP\_FULL\_DATABASE role, 7-530
- EXPLAIN PLAN statement, 7-512
- explicit data conversion, 2-37, 2-38
- expressions
  - CASE, 5-13
  - CAST, 5-7
  - compound, 5-4
  - computing with the DUAL table, 5-26
- CURSOR, 5-10
- DECODE, 5-12
- function, 5-5
- function, built-in, 5-5
- in SQL syntax, 5-1
- list of, 5-14
- object access, 5-11
- simple, 5-3
- type constructor, 5-6
- variable, 5-4
- extended rowids, 2-27
  - not directly available, 2-27
- EXTENT MANAGEMENT clause
  - for temporary tablespaces, 7-425
  - of CREATE TABLESPACE, 7-420, 7-422
- extents
  - allocating for partitions, 7-150
  - allocating for subpartitions, 7-150
  - allocating for tables, 7-150
  - restricting access by instances, 7-35
  - specifying maximum number for an  
object, 7-608
  - specifying number allocated upon object  
creation, 7-607
  - specifying the first for an object, 7-607
  - specifying the percentage of size increase, 7-607
  - specifying the second for an object, 7-607
- external datatypes
  - associated with host variables, 2-10
  - recognized by precompilers, 2-10
- external functions, 7-285, 7-354
- external LOBs, 2-23
- external procedures, 7-354
- external users, 7-366, 7-452

## F

---

- FAILED\_LOGIN\_ATTEMPTS parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- FAST\_START\_IO\_TARGET parameter
  - of ALTER SESSION, 7-87, 7-110
- FAST\_START\_PARALLEL\_ROLLBACK parameter
  - of ALTER SYSTEM, 7-110

- files
  - specifying as a redo log file group, 7-516
  - specifying as datafiles, 7-516
  - specifying as tempfiles, 7-516
- filespec clause, 7-516
  - of CREATE CONTROLFILE, 7-263
  - of CREATE DATABASE, 7-268
  - of CREATE LIBRARY, 7-316
  - of CREATE TABLESPACE, 7-419
  - of CREATE TEMPORARY TABLESPACE, 7-424
- FIPS compliance, B - 3
- FIPS flagging, 7-88
- FIRST\_ROWS hint, 2-64
- FIRST\_VALUE function, 4-34
- FIXED\_DATE parameter
  - of ALTER SYSTEM, 7-110
- FLAGGER parameter
  - of ALTER SESSION, 7-88
- FLOAT datatype, 2-19
  - ANSI, 2-28
  - DB2, 2-29
  - SQL/DS, 2-29
- floating-point numbers, 2-19
- FLOOR function, 4-35
- FLUSH SHARED POOL clause
  - of ALTER SYSTEM, 7-107
- FM format model modifier, 2-51
- FM number format element, 2-41
- FOR CATEGORY clause
  - of CREATE OUTLINE, 7-342
- FOR clause
  - of ANALYZE ... COMPUTE STATISTICS, 7-204
  - of ANALYZE ... ESTIMATE STATISTICS, 7-204
  - of CREATE INDEXTYPE, 7-310
  - of CREATE SYNONYM, 7-378
  - of EXPLAIN PLAN, 7-514
- FOR EACH ROW clause
  - of CREATE TRIGGER, 7-432
- FOR UPDATE clause
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-327
  - of SELECT, 7-572, 7-580
- FORCE ANY TRANSACTION system  
privilege, 7-529
- FORCE CLAUSE
  - of DROP OPERATOR, 7-489
- FORCE clause
  - of COMMIT, 7-231
  - of CREATE VIEW, 7-458
  - of DISASSOCIATE STATISTICS, 7-471
  - of DROP INDEX, 7-481
  - of DROP INDEXTYPE, 7-482
  - of DROP TYPE, 7-507
  - of REVOKE schema\_object\_privileges, 7-561
  - of ROLLBACK, 7-565
- FORCE PARALLEL DML clause
  - of ALTER SESSION, 7-84
- FORCE TRANSACTION system privilege, 7-529
- FOREIGN KEY clause
  - of constraint\_clause, 7-236, 7-240
- foreign key constraints, 7-240
- foreign tables
  - rowids of, 2-27
- format models, 2-38
  - changing the return format, 2-39
  - date, 2-45
  - date, changing, 2-45
  - date, default, 2-45
  - date, format elements, 2-45
  - date, maximum length, 2-45
  - modifiers, 2-51
  - number, 2-40
  - number, elements of, 2-41
  - specifying, 2-40
- formats
  - for dates and numbers. *See* format models
  - of return values from the database, 2-38
  - of values stored in the database, 2-38
- free lists
  - specifying for a table, partition, cluster, or  
index, 7-608
- FREELIST GROUPS parameter
  - of STORAGE clause, 7-608
- FREELISTS parameter
  - of STORAGE clause, 7-608
- FROM clause
  - of queries, 5-23
  - of REVOKE system\_privileges\_and\_roles, 7-557
  - of SELECT, 7-574
- FROM COLUMNS clause

- of DISASSOCIATE STATISTICS, 7-471
- FROM FUNCTIONS clause
  - of DISASSOCIATE STATISTICS, 7-471
- FROM INDEXES clause
  - of DISASSOCIATE STATISTICS, 7-471
- FROM INDEXTYPES clause
  - of DISASSOCIATE STATISTICS, 7-471
- FROM PACKAGES clause
  - of DISASSOCIATE STATISTICS, 7-471
- FROM PUBLIC clause
  - of REVOKE schema\_object\_privileges, 7-561
- FROM role clause
  - of REVOKE schema\_object\_privileges, 7-561
- FROM TYPES clause
  - of DISASSOCIATE STATISTICS, 7-471
- FROM user clause
  - of REVOKE schema\_object\_privileges, 7-561
- FULL hint, 2-65
- function expressions
  - built-in, 5-5
- function-based indexes, 7-295
  - and query rewrite, 7-93
  - creating, 7-298
  - disabling, 7-40, 7-116
  - enabling, 7-37, 7-40, 7-116
- functions
  - See also* SQL functions
  - 3GL, calling, 7-316
  - access to tables and packages, 7-442
  - associating statistics with, 7-212
  - avoiding run-time compilation, 7-28
  - calling, 7-226
  - changing the declaration of, 7-28
  - datatype of return value, 7-288
  - declaring
    - as a Java method, 7-289
    - as C functions, 7-289
  - defining an index on, 7-298
  - disassociating statistics types from, 7-478
  - examples, 7-289
  - executing, 7-226
    - from parallel query process, 7-288
  - expressions, 5-5
  - external, 7-285, 7-354
  - invalidating local objects dependent on, 7-478

- issuing COMMIT or ROLLBACK
  - statements, 7-84
- naming rules, 2-74
- privileges executed with, 7-188, 7-441
- recompiling invalid, 7-28
- re-creating, 7-286, 7-312
- removing from the database, 7-478
- schema executed in, 7-188, 7-441
- specification of
  - in CREATE TYPE, 7-442
- specifying schema and user privileges for, 7-288
- stored, 7-285
- storing return value of, 7-227
- synonyms for, 7-377
- user-defined, 4-118
- using a saved copy of, 7-288
- FUNCTIONS clause
  - of ASSOCIATE STATISTICS, 7-210, 7-212
- FX format model modifier, 2-51

## G

---

- G number format element, 2-41
- GC\_DEFER\_TIME parameter
  - of ALTER SYSTEM, 7-110
- general recovery clause
  - of ALTER DATABASE, 7-7, 7-12
- GLOBAL PARTITION BY RANGE clause
  - of CREATE INDEX, 7-302
- GLOBAL QUERY REWRITE system
  - privilege, 7-523, 7-524, 7-526
- GLOBAL TEMPORARY clause
  - of CREATE TABLE, 7-391
- global users, 7-366, 7-452
- GLOBAL\_NAMES parameter
  - of ALTER SESSION, 7-88
  - of ALTER SYSTEM, 7-110
- globally partitioned indexes, 7-302
- GRANT ANY PRIVILEGE system privilege, 7-529
- GRANT ANY ROLE system privilege, 7-525
- GRANT CONNECT THROUGH clause
  - of ALTER USER, 7-194, 7-195
- GRANT object\_privileges, 7-532
- GRANT system\_privileges\_and\_roles
  - statement, 7-519



GRAPHIC datatype (SQL/DS or DB2), 2-30  
greater than or equal to tests, 3-5  
greater than tests, 3-5  
GREATEST function, 4-36  
GROUP BY clause  
    of SELECT, 7-578  
    of SELECT and subqueries, 7-571  
group comparison conditions, 5-16  
GROUPING function, 4-36

## H

---

hash clusters  
    creating, 7-257  
    single-table, creating, 7-257  
    specifying hash function for, 7-257  
HASH hint, 2-65  
HASH IS clause  
    of CREATE CLUSTER, 7-257  
hash partition  
    adding, 7-160  
hash partitioning clause  
    of CREATE TABLE, 7-387, 7-402  
HASH\_AJ hint, 2-65  
HASH\_AREA\_SIZE parameter  
    of ALTER SESSION, 7-88  
HASH\_JOIN\_ENABLED parameter  
    of ALTER SESSION, 7-88  
HASH\_MULTIBLOCK\_IO\_COUNT parameter  
    of ALTER SESSION, 7-88  
    of ALTER SYSTEM, 7-111  
HASH\_SJ hint, 2-65  
HASHKEYS clause  
    of CREATE CLUSTER, 7-257  
HAVING condition  
    of GROUP BY clause, 7-579  
heap-organized tables  
    creating, 7-390  
hexadecimal value  
    returning, 2-43  
HEXTORAW function, 2-37, 4-37  
HH date format element, 2-46  
HH12 date format element, 2-46  
HH24 date format element, 2-46  
hierarchical queries, 2-58, 5-20, 7-577

    child nodes of, 2-58  
    child rows of, 5-20  
    illustrated, 2-59  
    leaf nodes of, 2-58  
    parent nodes of, 2-58  
    parent rows of, 5-20  
hierarchical query clause  
    of SELECT and subqueries, 7-571  
hierarchies  
    adding to a dimension, 7-27  
    dropping from a dimension, 7-27  
    of dimensions, defining, 7-279  
HIERARCHY clause  
    of ALTER DIMENSION, 7-25  
    of CREATE DIMENSION, 7-279  
high water mark  
    of clusters, 7-5  
    of indexes, 7-35  
    of tables, 7-150, 7-203  
hints, 5-19  
    in SQL statements, 2-63  
    passing to the optimizer, 7-616  
    syntax, 2-64  
HS\_ADMIN\_ROLE role, 7-530  
HS\_AUTOREGISTER parameter  
    of ALTER SYSTEM, 7-111

## I

---

I date format element, 2-46  
IDENTIFIED BY clause  
    of ALTER ROLE. *See* CREATE ROLE  
    of ALTER ROLE. *See* CREATE ROLE.  
    of CREATE ROLE, 7-366  
IDENTIFIED BY password clause  
    of CREATE DATABASE LINK, 7-275  
    of CREATE USER, 7-452  
    of SET ROLE, 7-601  
IDENTIFIED EXTERNALLY clause  
    of ALTER ROLE. *See* CREATE ROLE  
    of ALTER USER. *See* CREATE USER  
    of CREATE ROLE, 7-366  
    of CREATE USER, 7-452  
IDENTIFIED GLOBALLY clause  
    of ALTER ROLE. *See* CREATE ROLE

- of ALTER ROLE. *See* CREATE ROLE.
- of ALTER USER, 7-195
- of CREATE ROLE, 7-366
- of CREATE USER, 7-452
- of CREATE USERIDENTIFIED BY clause
  - of ALTER USER. *See* CREATE USER
- IDLE\_TIME parameter
  - of ALTER PROFILE, 7-69
  - of CREATE PROFILE, 7-361
- IMMEDIATE clause
  - of SET CONSTRAINTS, 7-598
- IMP\_FULL\_DATABASE role, 7-530
- implicit data conversion, 2-36, 2-38
- IN OUT parameter
  - of CREATE FUNCTION, 7-287
  - of CREATE PROCEDURE, 7-356
- IN parameter
  - of CREATE function, 7-287
  - of CREATE PROCEDURE, 7-356
- INCLUDING clause
  - of ALTER TABLE, 7-154
- INCLUDING CONTENTS clause
  - of DROP TABLESPACE, 7-504
- INCLUDING NEW VALUES clause
  - of ALTER MATERIALIZED VIEW LOG, 7-61
  - of CREATE MATERIALIZED VIEW LOG /  
SNAPSHOT LOG, 7-337
- INCLUDING TABLES clause
  - of DROP CLUSTER, 7-473
- incomplete object types, 7-439
  - creating, 7-437, 7-439
- INCREMENT BY clause
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE
  - of CREATE SEQUENCE, 7-373
- INDEX clause
  - of CREATE CLUSTER, 7-257
- INDEX hint, 2-65
- INDEX object privilege, 7-535
- index partitions
  - changing physical attributes of, 7-36
  - deallocating unused space from, 7-35
  - dropping, 7-41
  - marking UNUSABLE, 7-157
  - modifying the real characteristics of, 7-41
  - rebuilding, 7-37
    - unusable, 7-157
  - renaming, 7-41
    - specifying tablespace for, 7-38
  - splitting, 7-41
- index subpartitions
  - allocating extents for, 7-42
  - changing physical attributes of, 7-36
  - deallocating unused space from, 7-35, 7-42
  - marking UNUSABLE, 7-42
  - rebuilding, 7-37
  - renaming, 7-41
    - specifying tablespace for, 7-38
- INDEX\_ASC hint, 2-65
- INDEX\_COMBINE hint, 2-65
- INDEX\_DESC hint, 2-65
- INDEX\_FFS hint, 2-65
- indexed clusters
  - creating, 7-257
- indexes
  - allocating new extents for, 7-35
  - application-specific, 7-309
  - ascending, 7-299
  - based on indextypes, 7-304
  - bitmap, 7-297
  - changing attributes of, 7-36
  - collecting statistics on, 7-202
  - on composite-partitioned tables, 7-303
  - creating, 7-295
  - creating as cluster indexes, 7-297
  - deallocating unused space from, 7-35
  - descending, 7-299
    - and query rewrite, 7-299
    - as function-based indexes, 7-299
  - disassociating statistics types from, 7-480
  - domain, 7-295, 7-304, 7-309
  - dropping index partitions of, 7-480
  - examples, 7-305
  - function-based, 7-295
    - creating, 7-298
  - globally partitioned, 7-302
  - granting
    - system privileges on, 7-522
  - on hash-partitioned tables, 7-303
  - key compression of, 7-39, 7-300
  - locally partitioned, 7-302

- logging attributes of, 7-301
- logging rebuild operations on, 7-39
- marking as UNUSABLE, 7-40
- merging contents of index blocks, 7-40
- online, 7-301
- parallel queries and DML on, 7-36
- parallelizing creation of, 7-304
- partitioned, 2-70, 7-295
  - user-defined, 7-302
- partitions
  - adding new, 7-41
  - marking UNUSABLE, 7-41
- physical attributes of, 7-300
- on range-partitioned tables, 7-302
- rebuilding, 7-37
- rebuilding while online, 7-39
- removing from the database, 7-480
- renaming, 7-40
- reverse, 7-38, 7-300
- specifying tablespace for, 7-38
- statistics on, 7-301
- statistics on rebuild, 7-39
- storage characteristics of, 7-300, 7-605
- tablespace containing, 7-300
- unique, 7-296
- unsorted, 7-300
- validating structure of, 7-206

INDEXES clause

- of ASSOCIATE STATISTICS, 7-210, 7-212

index-organized table clause

- of CREATE TABLE, 7-384, 7-396

index-organized tables

- creating, 7-390, 7-396
- modifying, 7-153
- rebuilding, 7-144
- reserving space in the index block, 7-154
- rowids of, 2-27

INDEXTYPE clause

- of CREATE INDEX, 7-304

indextypes

- associating statistics with, 7-212
- creating, 7-309
- disassociating from statistics types, 7-482
- drop routines, invoking, 7-480
- granting
  - system privileges on, 7-522
- indexes based on, 7-304
- instances of, 7-295
- removing from the database, 7-482

INDEXTYPES clause

- of ASSOCIATE STATISTICS, 7-210, 7-212

in-doubt transactions

- forcing, 7-231
- forcing commit of, 7-231
- forcing rollback, 7-565
- forcing rollback of, 7-565
- rolling back, 7-564

inequality test, 3-5

INITCAP function, 4-38

INITIAL parameter

- of STORAGE clause, 7-607

INITIALLY DEFERRED clause

- of constraint\_clause, 7-243

INITIALLY IMMEDIATE clause

- of constraint\_clause, 7-243

INITRANS parameter

- of ALTER CLUSTER, 7-4
- of ALTER INDEX, 7-31, 7-36
- of ALTER MATERIALIZED VIEW, 7-49
- of ALTER MATERIALIZED VIEW LOG, 7-58
- of CREATE CLUSTER, 7-256
- of CREATE INDEX. *See* CREATE TABLE
- of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG. *See* CREATE TABLE
- of CREATE MATERIALIZED VIEW/  
SNAPSHOT. *See* CREATE TABLE
- of CREATE TABLE, 7-394

INSERT ANY TABLE system privilege, 7-527

INSERT object privilege, 7-535

INSERT statement, 7-539

- triggers on, 7-429

INSTANCE parameter

- of ALTER SESSION, 7-88

instances

- global name resolution for, 7-110
- setting parameters for, 7-108

INSTEAD OF clause

- of CREATE TRIGGER, 7-429

INSTEAD OF triggers, 7-429

INSTR function, 4-38

INSTRB function, 4-39  
INT datatype (ANSI), 2-28  
INTEGER datatype  
  ANSI, 2-28  
  DB2, 2-29  
  SQL/DS, 2-29  
integers  
  generating unique, 7-371  
  in SQL syntax, 2-3  
  precision of, 2-3  
  specifying, 2-17  
  syntax of, 2-3  
integrity constraints. *See* constraints  
internal LOBs, 2-23  
International Standards Organization. *See* ISO  
INTERSECT set operator, 3-13, 7-579  
INTO clause  
  of EXPLAIN PLAN, 7-513  
  of INSERT, 7-540  
INTO host\_variable clause  
  of CALL, 7-227  
invoker rights clause  
  of ALTER JAVA, 7-46  
  of CREATE FUNCTION, 7-288  
  of CREATE JAVA, 7-313  
  of CREATE PACKAGE, 7-344  
  of CREATE PROCEDURE, 7-353  
  of CREATE TYPE, 7-188, 7-441  
IS NOT NULL operator, 3-7  
IS NULL operator, 3-7  
ISO, B - 1  
  standards, xv, 1-2, B - 1  
ISOLATION LEVEL READ COMMITTED clause  
  of SET TRANSACTION, 7-603  
ISOLATION LEVEL SERIALIZABLE clause  
  of SET TRANSACTION, 7-603  
ISOLATION\_LEVEL parameter  
  of ALTER SESSION, 7-89  
IW date format element, 2-46  
IY date format element, 2-46  
IYY date format element, 2-46  
IYYY date format element, 2-46

## J

---

J date format element, 2-46  
Java class schema object  
  creating, 7-312, 7-313  
  dropping, 7-483  
  resolving, 7-45, 7-312  
JAVA clause  
  of CREATE TYPE, 7-442  
  of CREATE TYPE BODY, 7-450  
Java method  
  mapping to an object type, 7-442  
Java resource schema object  
  creating, 7-312, 7-313  
  dropping, 7-483  
Java schema object  
  name resolution of, 7-314  
Java source schema object  
  compiling, 7-45, 7-312  
  creating, 7-312, 7-313  
  dropping, 7-483  
JOB\_QUEUE\_PROCESSES parameter  
  of ALTER SYSTEM, 7-111  
JOIN KEY clause  
  of ALTER DIMENSION, 7-26  
  of CREATE DIMENSION, 7-280  
join views  
  modifying, 7-466, 7-541, 7-618  
joins, 5-22  
  conditions  
    defining, 5-22  
  equijoins, 5-22  
  outer, 5-23  
    restrictions, 5-23  
  self, 5-23  
  without join conditions, 5-23  
Julian day, specifying, 2-47

## K

---

key compression, 7-39, 7-300, 7-397  
  disabling, 7-39, 7-300  
  of index rebuild, 7-144  
  of indexes, 7-39, 7-300  
    disabling, 7-39

- of index-organized tables, 7-397
- keywords, 2-73
  - in syntax diagrams, xxi
  - optional, A-2
  - required, A-1
- KILL SESSION clause
  - of ALTER SYSTEM, 7-107

## L

---

- L number format element, 2-41
- LAG function, 4-40
- LANGUAGE clause
  - of CREATE FUNCTION, 7-289
  - of CREATE PROCEDURE, 7-357
  - of CREATE TYPE, 7-442
  - of CREATE TYPE BODY, 7-450
- large objects. *See* LOBs
- LAST\_DAY function, 4-41
- LAST\_VALUE function, 4-42
- LEAD function, 4-44
- LEAST function, 4-45
- LENGTH function, 4-45
- LENGTHB function, 4-46
- less than tests, 3-5
- LEVEL clause
  - of ALTER DIMENSION, 7-25
  - of CREATE DIMENSION, 7-279
- LEVEL pseudocolumn, 2-58, 7-577
  - and hierarchical queries, 2-58
- levels
  - adding to a dimension, 7-27
  - dropping from a dimension, 7-27
  - of dimensions, defining, 7-279
- libraries
  - creating, 7-316
  - granting
    - system privileges on, 7-523
  - re-creating, 7-316
  - removing from the database, 7-484
- library units. *See* Java schema objects
- LICENSE\_MAX\_SESSIONS parameter
  - of ALTER SYSTEM, 7-111
- LICENSE\_MAX\_USERS parameter
  - of ALTER SYSTEM, 7-112
- LICENSE\_SESSIONS\_WARNING parameter
  - of ALTER SYSTEM, 7-112
- LIKE conditions, 5-18
- LIKE operator, 3-8
- linear regression functions, 4-71
- LIST CHAINED ROWS clause
  - of ANALYZE, 7-207
- literals
  - in SQL statements and functions, 2-2
  - in SQL syntax, 2-2
- LN function, 4-46
- LOB datatypes, 2-23
- LOB index clause
  - of ALTER TABLE, 7-142
  - of CREATE TABLE, 7-399
- LOB storage clause
  - for partitions, 7-142
  - of ALTER MATERIALIZED VIEW, 7-47, 7-51
  - of ALTER TABLE, 7-141
  - of CREATE MATERIALIZED VIEW /  
SNAPSHOT, 7-319
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-322
  - of CREATE TABLE, 7-385, 7-398
- LOBs
  - attributes
    - initializing, 2-24
  - columns
    - difference from LONG and LONG  
RAW, 2-23
    - populating, 2-24
  - external, 2-23
  - indexes for, 7-399
  - internal, 2-23
  - locators, 2-23, 7-399
  - logging attribute of, 7-395
  - modifying physical attributes of, 7-146
  - number of bytes manipulated in, 7-399
  - specifying directories for, 7-282
  - storage
    - characteristics, 7-394, 7-398
    - in-line, 7-398
    - outside of row, 7-398
  - tablespace for
    - defining, 7-394

- LOCAL clause
  - of CREATE INDEX, 7-302
- local users, 7-366, 7-452
- locally managed tablespaces
  - storage characteristics, 7-606
- locally partitioned indexes, 7-302
- LOCK ANY TABLE system privilege, 7-527
- LOCK TABLE statement, 7-547
- locking
  - automatic
    - overriding, 7-547
- locks. *See* table locks
- LOG function, 4-47
- LOG\_ARCHIVE\_DEST parameter
  - of ALTER SYSTEM, 7-112
- LOG\_ARCHIVE\_DEST\_n parameter
  - of ALTER SESSION, 7-89, 7-112
  - of ALTER SYSTEM, 7-112
- LOG\_ARCHIVE\_DEST\_STATE\_n parameter
  - of ALTER SESSION, 7-89
  - of ALTER SYSTEM, 7-113
- LOG\_ARCHIVE\_DUPLEX\_DEST parameter
  - of ALTER SYSTEM, 7-113
- LOG\_ARCHIVE\_MAX\_PROCESSES parameter
  - of ALTER SYSTEM, 7-113
- LOG\_ARCHIVE\_MIN\_SUCCEED\_DEST parameter
  - of ALTER SESSION, 7-89
  - of ALTER SYSTEM, 7-113
- LOG\_ARCHIVE\_TRACE parameter
  - of ALTER SYSTEM, 7-113
- LOG\_CHECKPOINT\_INTERVAL parameter
  - of ALTER SYSTEM, 7-113
- LOG\_CHECKPOINT\_TIMEOUT parameter
  - of ALTER SYSTEM, 7-114
- LOGFILE clause
  - of CREATE CONTROLFILE, 7-265
  - OF CREATE DATABASE, 7-269
- logfile clauses
  - of ALTER DATABASE, 7-9
- LOGFILE GROUP clause
  - of CREATE CONTROLFILE, 7-265
- logging
  - and redo log size, 7-396
  - specifying minimal, 7-395
- LOGGING clause
  - of ALTER INDEX, 7-36
  - of ALTER MATERIALIZED VIEW, 7-50
  - of ALTER MATERIALIZED VIEW LOG, 7-60
  - of ALTER TABLE, 7-152
  - of ALTER TABLESPACE, 7-178
  - of CREATE INDEX, 7-301
  - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-322
  - of CREATE TABLE, 7-395
  - of CREATE TABLESPACE, 7-421
- logical operators, 3-11
- LOGICAL\_READS\_PER\_CALL parameter
  - of ALTER PROFILE, 7-69
  - of CREATE PROFILE, 7-362
- LOGICAL\_READS\_PER\_SESSION parameter
  - of ALTER PROFILE, 7-69
  - of ALTER RESOURCE COST, 7-74
  - of CREATE PROFILE, 7-362
- LOGOFF
  - triggers on, 7-431
- LOGOFF event
  - triggers on, 7-431
- LOGON
  - triggers on, 7-431
- LOGON event
  - triggers on, 7-431
- LONG columns
  - converting to LOB columns, 2-19, 2-22
  - restrictions on, 2-20
  - to store text strings, 2-19
  - to store view definitions, 2-19
  - where referenced from, 2-20
- LONG datatype, 2-19
  - in triggers, 2-20
- LONG RAW
  - data
    - converting from CHAR data, 2-23
  - datatype, 2-22
- LONG VARCHAR datatype
  - DB2, 2-29
  - SQL/DS, 2-29
- LONG VARGRAPHIC datatype (SQL/DS or  
DB2), 2-30

LOWER function, 4-47  
LPAD function, 4-48  
LTRIM function, 4-48

## M

---

MAKE\_REF function, 4-49  
MANAGE TABLESPACE system privilege, 7-527  
managed recovery  
  of database, 7-8  
MANAGED STANDBY RECOVERY clause  
  of ALTER DATABASE, 7-14  
MAP MEMBER clause  
  of ALTER TYPE, 7-189  
  of CREATE TYPE, 7-443, 7-449  
MAP methods  
  specifying, 7-189  
master databases, 7-320  
master tables, 7-320  
materialized join views, 7-335  
materialized view logs, 7-334  
  creating, 7-334  
  excluding new values from, 7-61  
  logging changes to, 7-60  
  parallelizing creation of, 7-336  
  partition attributes  
    changing, 7-60  
  partitioned, 7-336  
  physical attributes  
    changing, 7-60  
    specifying, 7-336  
  removing from the database, 7-487  
  required for fast refresh, 7-334  
  saving new values in, 7-61  
  saving old values in, 7-337  
  storage characteristics  
    specifying, 7-336  
materialized views  
  allowing update of, 7-327  
  complete refresh, 7-52, 7-325  
  constraints on, 7-244  
  creating, 7-320  
  creating comments about, 7-228  
  for data warehousing, 7-320  
  degree of parallelism, 7-51, 7-60

  during creation, 7-323  
  detail table of, dropping, 7-486  
  enabling and disabling query rewrite, 7-327  
  examples, 7-329, 7-337  
  fast refresh, 7-52, 7-324, 7-325  
  forced refresh, 7-53  
  from existing tables, 7-323  
  granting  
    system privileges on, 7-523  
  index characteristics  
    changing, 7-51  
  indexes that maintain, 7-324  
  join, 7-335  
  LOB storage characteristics of, 7-51  
  logging changes to, 7-50  
  partitions of, 7-51  
  physical and storage attributes  
    changing, 7-50  
  physical attributes of, 7-322  
  primary key, 7-326  
    recording values in master table, 7-61  
  query rewrite  
    eligibility for, 7-244  
    enabling and disabling, 7-54  
  re-creating during refresh, 7-52  
  refresh mode  
    changing, 7-52  
  refresh time  
    changing, 7-52  
  refreshing after DML on master table, 7-53,  
    7-325  
  refreshing on next COMMIT, 7-53, 7-325  
  removing from the database, 7-485  
  for replication, 7-320  
  retrieving data from, 7-572  
  revalidating, 7-54  
  rowid, 7-326  
  rowid values  
    recording in master table, 7-61  
  saving blocks in a cache, 7-51  
  storage characteristics of, 7-322  
  subquery, 7-328  
  synonyms for, 7-377  
  when to populate, 7-323  
MAX function, 4-50

- MAX\_DUMP\_FILE\_SIZE parameter
  - of ALTER SESSION, 7-90
  - of ALTER SYSTEM, 7-114
- MAXDATAFILES parameter
  - of CREATE CONTROLFILE, 7-265
  - OF CREATE DATABASE, 7-270
- MAXEXTENTS parameter
  - of STORAGE clause, 7-608
- MAXINSTANCES parameter
  - of CREATE CONTROLFILE, 7-266
  - OF CREATE DATABASE, 7-270
- MAXLOGFILES parameter
  - of CREATE CONTROLFILE, 7-265
  - OF CREATE DATABASE, 7-270
- MAXLOGHISTORY parameter
  - of CREATE CONTROLFILE, 7-265
  - OF CREATE DATABASE, 7-270
- MAXLOGMEMBERS parameter
  - of CREATE CONTROLFILE, 7-265
  - OF CREATE DATABASE, 7-270
- MAXSIZE clause
  - of ALTER DATABASE, 7-10
  - of CREATE DATABASE, 7-268
  - of CREATE TABLESPACE, 7-420
  - of CREATE TEMPORARY TABLESPACE, 7-424
- MAXTRANS parameter
  - of ALTER CLUSTER, 7-4
  - of ALTER INDEX, 7-31, 7-36
  - of ALTER MATERIALIZED VIEW, 7-49
  - of ALTER MATERIALIZED VIEW LOG, 7-58
  - of CREATE CLUSTER, 7-256
  - of CREATE INDEX. *See* CREATE TABLE
  - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG. *See* CREATE TABLE
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT. *See* CREATE TABLE
  - of CREATE TABLE, 7-394
- MAXVALUE clause
  - of CREATE SEQUENCE, 7-373
- MAXVALUE parameter
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE
- media recovery
  - disabling, 7-17
  - of database, 7-12
  - of datafiles, 7-12
  - of standby database, 7-12
  - of tablespaces, 7-12
  - parallelizing, 7-15
  - restrictions, 7-13
  - sustained standby recovery, 7-14
- MEMBER clause
  - of ALTER TYPE, 7-188
  - of CREATE TYPE, 7-441
  - of CREATE TYPE BODY, 7-448
- membership conditions, 5-17
- MERGE hint, 2-67
- MERGE PARTITIONS clause
  - of ALTER TABLE, 7-162
- MERGE\_AJ hint, 2-65
- MERGE\_SJ hint, 2-65
- MI date format element, 2-46
- MI number format element, 2-41
- migrated rows
  - listing, 7-207
- MIN function, 4-52
- MINEXTENTS parameter
  - of STORAGE clause, 7-607
- MINIMIZE RECORDS PER BLOCK clause
  - of ALTER TABLE, 7-153
- MINIMUM EXTENT clause
  - of ALTER TABLESPACE, 7-180
  - of CREATE TABLESPACE, 7-421
- MINUS set operator, 3-13, 7-579
- MINVALUE
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE
- MINVALUE clause
  - of CREATE SEQUENCE, 7-373
- MM date format element, 2-46
- MOD function, 4-53
- MODE clause
  - of LOCK TABLE, 7-548
- MODIFY clause
  - of ALTER TABLE, 7-143
- MODIFY CONSTRAINT clause
  - of ALTER TABLE, 7-144
- MODIFY DEFAULT ATTRIBUTES clause
  - of ALTER INDEX, 7-33, 7-41
  - of ALTER TABLE, 7-155
- MODIFY LOB clause
  - of ALTER TABLE, 7-146



- MODIFY LOB storage clause
  - of ALTER MATERIALIZED VIEW, 7-48, 7-51
  - of ALTER TABLE, 7-146
- MODIFY NESTED TABLE clause
  - of ALTER TABLE, 7-145
- MODIFY PARTITION clause
  - of ALTER INDEX, 7-33, 7-41
  - of ALTER MATERIALIZED VIEW, 7-51
  - of ALTER TABLE, 7-156
- MODIFY SUBPARTITION clause
  - of ALTER INDEX, 7-34, 7-42
  - of ALTER TABLE, 7-157
- MODIFY VARRAY clause
  - of ALTER TABLE, 7-146
- modifying space for each cluster key, 7-4
- MON date format element, 2-46, 2-48
- MONITORING clause
  - of ALTER TABLE, 7-152
  - of CREATE TABLE, 7-408
- MONTH date format element, 2-46, 2-48
- MONTHS\_BETWEEN function, 4-54
- MOUNT clause
  - of ALTER DATABASE, 7-11
- MOVE clause
  - of ALTER TABLE, 7-144
- MOVE ONLINE clause
  - of ALTER TABLE, 7-144
- MOVE PARTITION clause
  - of ALTER TABLE, 7-157
- MOVE SUBPARTITION clause
  - of ALTER TABLE, 7-158
- MTS\_DISPATCHERS parameter
  - of ALTER SYSTEM, 7-114
- MTS\_SERVERS parameter
  - of ALTER SYSTEM, 7-115
- multi-threaded server
  - system parameters, 7-114

## N

---

- NAMED clause
  - of CREATE JAVA, 7-313
- namespaces
  - and object naming rules, 2-73
  - for nonschema objects, 2-74

- for schema objects, 2-73
- NATIONAL CHAR datatype (ANSI), 2-28
- NATIONAL CHAR VARYING datatype (ANSI), 2-28
- NATIONAL CHARACTER datatype (ANSI), 2-28
- national character set
  - fixed vs. variable width, 2-15, 2-16
  - multibyte character data, 2-26
  - multibyte character sets, 2-15, 2-16
  - variable-length strings, 2-16
- NATIONAL CHARACTER SET clause
  - of CREATE DATABASE, 7-271
- NATIONAL CHARACTER SET parameter
  - of ALTER DATABASE, 7-16
- NATIONAL CHARACTER VARYING datatype
  - ANSI, 2-28
- national language support. *See* NLS
- NCHAR datatype, 2-15
  - ANSI, 2-28
- NCHAR VARYING datatype (ANSI), 2-28
- NCLOB datatype, 2-26
  - transactional support of, 2-26
- negative scale, 2-18
- NESTED TABLE clause
  - of ALTER TABLE, 7-142
  - of CREATE TABLE, 7-386, 7-400
- nested table types, 2-31
  - compared with varrays, 2-35
  - comparison rules, 2-35
  - creating, 7-439
  - dropping the body of, 7-508
  - dropping the specification of, 7-506
  - modifying, 7-145
- nested tables
  - changing returned value, 7-145
  - creating, 7-444
  - defining as index-organized tables, 7-142
  - storage characteristics of, 7-142, 7-400
- NEW\_TIME function, 4-55
- NEXT clause
  - of ALTER MATERIALIZED VIEW...REFRESH, 7-53
- NEXT parameter
  - of STORAGE clause, 7-607
- NEXT\_DAY function, 4-55

NEXTVAL pseudocolumn, 2-56, 7-372  
 NLS parameters  
 NLS\_CALENDAR parameter  
   of ALTER SESSION, 7-90  
 NLS\_CHARSET\_DECL\_LEN function, 4-56  
 NLS\_CHARSET\_ID function, 4-56  
 NLS\_CHARSET\_NAME function, 4-57  
 NLS\_COMP parameter  
   of ALTER SESSION, 7-90  
 NLS\_CURRENCY parameter  
   of ALTER SESSION, 7-90  
 NLS\_DATE\_FORMAT parameter  
   of ALTER SESSION, 7-90  
 NLS\_DATE\_LANGUAGE parameter, 2-48  
   of ALTER SESSION, 7-90  
 NLS\_INITCAP function, 4-58  
 NLS\_ISO\_CURRENCY parameter  
   of ALTER SESSION, 7-90  
 NLS\_LANGUAGE parameter, 2-48, 5-21  
   of ALTER SESSION, 7-90  
 NLS\_LOWER function, 4-59  
 NLS\_NUMERIC\_CHARACTERS parameter  
   of ALTER SESSION, 7-91  
 NLS\_SORT parameter, 5-21  
   of ALTER SESSION, 7-91  
 NLS\_TERRITORY parameter, 2-48  
   of ALTER SESSION, 7-91  
 NLS\_UNION\_CURRENCY parameter  
   of ALTER SESSION, 7-90  
 NLS\_UPPER function, 4-60  
 NLSSORT function, 4-59  
 NO\_EXPAND hint, 2-65  
 NO\_INDEX hint, 2-65  
 NO\_MERGE hint, 2-67  
 NO\_PUSH\_JOIN\_PRED hint, 2-67  
 NOAPPEND hint, 2-67  
 NOARCHIVELOG clause  
   of ALTER DATABASE, 7-9, 7-18  
   of CREATE CONTROLFILE, 7-266  
   OF CREATE DATABASE, 7-270  
 NOAUDIT schema\_objects statement, 7-552  
 NOAUDIT sql\_statements statement, 7-550  
 NOCACHE clause  
   of ALTER MATERIALIZED VIEW, 7-51  
   of ALTER MATERIALIZED VIEW LOG, 7-61  
   of ALTER SEQUENCE. *See* CREATE SEQUENCE  
   of ALTER TABLE, 7-151  
   of CREATE CLUSTER, 7-258  
   of CREATE MATERIALIZED VIEW LOG/  
     SNAPSHOT LOG, 7-336  
   of CREATE MATERIALIZED VIEW/  
     SNAPSHOT, 7-323  
   of CREATE SEQUENCE, 7-374  
   of CREATE TABLE, 7-408  
 NOCACHE hint, 2-67  
 NOCOMPRESS clause  
   of ALTER TABLE, 7-144  
   of CREATE INDEX, 7-300  
   of CREATE TABLE, 7-397  
 NOCOPY clause  
   of CREATE FUNCTION, 7-287  
   of CREATE PROCEDURE, 7-356  
 NOCYCLE clause  
   of ALTER SEQUENCE. *See* CREATE  
     SEQUENCE, 7-81  
   of CREATE SEQUENCE, 7-373  
 NOFORCE clause  
   of CREATE JAVA, 7-313  
   of CREATE VIEW, 7-458  
 NOLOGGING clause  
   of ALTER INDEX, 7-36  
   of ALTER MATERIALIZED VIEW, 7-50  
   of ALTER MATERIALIZED VIEW LOG, 7-60  
   of ALTER TABLE, 7-152  
   of ALTER TABLESPACE, 7-178  
   of CREATE INDEX, 7-301  
   of CREATE MATERIALIZED VIEW LOG/  
     SNAPSHOT LOG, 7-336  
   of CREATE MATERIALIZED VIEW/  
     SNAPSHOT, 7-322  
   of CREATE TABLE, 7-395  
   of CREATE TABLESPACE, 7-421  
 NOMAXVALUE clause  
   of CREATE SEQUENCE, 7-373  
 NOMAXVALUE parameter  
   of ALTER SEQUENCE. *See* CREATE SEQUENCE  
 NOMINIMIZE RECORDS PER BLOCK clause  
   of ALTER TABLE, 7-153  
 NOMINVALUE clause  
   of ALTER SEQUENCE. *See* CREATE

- SEQUENCE, 7-81
  - of CREATE SEQUENCE, 7-373
- NOMONITORING clause
  - of ALTER TABLE, 7-152
  - of CREATE TABLE, 7-409
- NONE clause
  - of SET ROLE, 7-601
- nonequivalency tests, 3-6
- nonschema objects
  - list of, 2-69
  - namespaces, 2-74
- NOORDER clause
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE
  - of CREATE SEQUENCE, 7-374
- NOPARALLEL clause
  - of CREATE CLUSTER, 7-258
  - of CREATE INDEX, 7-304
  - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-323
  - of CREATE TABLE, 7-405
- NOPARALLEL hint, 2-67
- NOPARALLEL\_INDEX hint, 2-67
- NORELY clause
  - of constraint\_clause, 7-244
- NORESETLOGS clause
  - of CREATE CONTROLFILE, 7-265
- NOREWRITE hint, 2-65
- NOSORT clause
  - of ALTER INDEX, 7-300
  - of constraint\_clause, 7-244
- NOT DEFERRABLE clause
  - of constraint\_clause, 7-243
- NOT IDENTIFIED clause
  - of ALTER ROLE. *See* CREATE ROLE
  - of CREATE ROLE, 7-366
- NOT IN operator, 3-7
- NOT NULL clause
  - of constraint\_clause, 7-239
  - of CREATE TABLE, 7-393
- NOT NULL constraints, 7-239
- not null constraints, 7-239
- NOT operator, 3-11
- NOWAIT clause

- of LOCK TABLE, 7-548
- NTILE function, 4-61
- null, 2-54
  - difference from zero, 2-54
  - in conditions, 2-55
    - table of, 2-55
  - in functions, 2-54
  - with comparison operators, 2-55
- NULL clause
  - of constraint\_clause, 7-239
- NULL conditions, 5-18
- NUMBER datatype, 2-17
  - converting to VARCHAR2, 2-40
  - precision, 2-17
  - scale, 2-17
- number format models, 2-40
- number functions, 4-4
- numbers
  - comparison rules, 2-32
  - floating-point, 2-17, 2-19
  - in SQL syntax, 2-4
  - precision of, 2-4
  - rounding, 2-18
  - spelling out, 2-51
  - syntax of, 2-4
- NUMERIC datatype (ANSI), 2-28
- NUMTODSINTERVAL function, 4-62
- NUMTOYMINTERVAL function, 4-63
- NVARCHAR2 datatype, 2-16
- NVL function, 4-64
- NVL2 function, 4-65

## O

---

- object access expressions, 5-11
- object cache, 7-91, 7-115
- OBJECT IDENTIFIER clause
  - of CREATE TABLE, 7-401
- object identifiers
  - contained in REFs, 2-31
  - of object views, 7-459
  - primary key, 7-401
  - specifying, 7-401
  - specifying an index on, 7-402
  - system-generated, 7-401

- object privileges
  - granting, 7-365
  - multiple, 7-369
  - on specific columns, 7-533
  - to a role, 7-532
  - to a user, 7-532
  - to PUCLIC, 7-533
- on a database object
  - revoking, 7-561
- on a directory
  - revoking, 7-560
- revoking
  - all, from a user, 7-560
  - from a role, 7-559, 7-560
  - from a user, 7-559, 7-560
  - from all users, 7-561
  - from PUBLIC, 7-560
- object reference functions, 4-12
- object tables
  - adding rows to, 7-540
  - creating, 7-381
- object type bodies
  - creating, 7-448
  - re-creating, 7-448
  - SQL examples, 7-450
- object type tables
  - creating, 7-391
- object type values
  - comparing, 7-443, 7-449
- object types, 2-30
  - adding new member subprograms, 7-188
  - associating functions or procedures, 7-188
  - attributes, 2-82
  - comparison rules, 2-35
    - MAP function, 2-35
    - ORDER function, 2-35
  - compiling the specification and body, 7-187
  - components of, 2-30
  - creating, 7-437, 7-439
  - defining member methods of, 7-448
  - disassociating statistics types from, 7-506
  - dropping the body of, 7-508
  - dropping the specification of, 7-506
  - function subprogram
    - declaring, 7-449
    - function subprogram of
      - specifying, 7-442
  - function subprograms of, 7-441, 7-448
  - granting system privileges on, 7-528
  - incomplete, 7-437, 7-439
  - methods, 2-82
  - nested table, 7-439
  - procedure subprogram
    - declaring, 7-449
  - procedure subprogram of
    - specifying, 7-442
  - procedure subprograms of, 7-441, 7-448
  - SQL examples, 7-444
  - statistics types, 7-211
  - user-defined
    - creating, 7-440
  - varrays, 7-439
- object views
  - adding rows to the base table of, 7-540
  - defining, 7-456
- OBJECT\_CACHE\_MAX\_SIZE\_PERCENT parameter
  - of ALTER SESSION, 7-91
  - of ALTER SYSTEM, 7-115
- OBJECT\_CACHE\_OPTIMAL\_SIZE parameter
  - of ALTER SESSION, 7-91
  - of ALTER SYSTEM, 7-115
- objects. *See* object types or database objects
- OF clause
  - of CREATE VIEW, 7-459
- OF object\_type clause
  - of CREATE TABLE, 7-391
- OFFLINE clause
  - of ALTER ROLLBACK SEGMENT, 7-79
  - of ALTER TABLESPACE, 7-180
  - of CREATE TABLESPACE, 7-422
- OIDINDEX clause
  - of CREATE TABLE, 7-402
- OIDs. *See* object identifiers
- ON clause
  - of CREATE OUTLINE, 7-342
- ON COMMIT clause
  - of CREATE TABLE, 7-401
- ON DATABASE clause
  - of CREATE TRIGGER, 7-432
- ON DEFAULT clause

- of AUDIT schema\_objects, 7-222
- of NOAUDIT schema\_objects, 7-552
- ON DELETE CASCADE clause
  - of constraint\_clause, 7-241
- ON DELETE SET NULL clause
  - of constraint\_clause, 7-241
- ON DIRECTORY clause
  - of AUDIT schema\_objects, 7-222
  - of GRANT object\_privileges, 7-533
  - of NOAUDIT schema\_objects, 7-552
  - of REVOKE schema\_object\_privileges, 7-560
- ON JAVA RESOURCE clause
  - of GRANT object\_privileges, 7-533
- ON JAVA SOURCE clause
  - of GRANT object\_privileges, 7-533
- ON NESTED TABLE clause
  - of CREATE TRIGGER, 7-432
- ON object clause
  - of GRANT object\_privileges, 7-533
  - of NOAUDIT schema\_objects, 7-552
  - of REVOKE schema\_object\_privileges, 7-561
- ON PREBUILT TABLE clause
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-323
- ON SCHEMA clause
  - of CREATE TRIGGER, 7-432
- ONLINE clause
  - of ALTER ROLLBACK SEGMENT, 7-78
  - of ALTER TABLESPACE, 7-180
  - of CREATE INDEX, 7-301
  - of CREATE TABLESPACE, 7-422
- online indexes, 7-301
  - rebuilding, 7-144
- online redo logs
  - reinitializing, 7-19
- OPEN NORESETLOGS clause
  - of ALTER DATABASE, 7-12
- OPEN READ ONLY clause
  - of ALTER DATABASE, 7-12
- OPEN READ WRITE clause
  - of ALTER DATABASE, 7-12
- OPEN RESETLOGS clause
  - of ALTER DATABASE, 7-12
- operands, 3-1
- operators, 3-1
  - arithmetic, 3-3
  - binary, 3-1
  - comparison, 3-5
  - concatenation, 3-3
  - granting
    - system privileges on, 7-524
  - logical, 3-11
  - precedence, 3-2
  - set, 3-12, 7-579
  - unary, 3-1
  - user-defined, 3-16
    - binding to a function, 7-340
    - creating, 7-340
    - dropping, 7-489
    - function providing implementation, 7-341
    - how bindings are implemented, 7-341
    - implementation type, 7-341
    - return type of binding, 7-340
- OPTIMAL parameter
  - of STORAGE clause, 7-368, 7-608
- optimization
  - setting session parameters, 7-92
- OPTIMIZER\_INDEX\_CACHING parameter
  - of ALTER SESSION, 7-92
- OPTIMIZER\_INDEX\_COST\_ADJ parameter
  - of ALTER SESSION, 7-92
- OPTIMIZER\_MAX\_PERMUTATIONS parameter
  - of ALTER SESSION, 7-92, 7-115
- OPTIMIZER\_MODE parameter
  - of ALTER SESSION, 7-92
- OPTIMIZER\_PERCENT\_PARALLEL parameter
  - of ALTER SESSION, 7-92
- OR operator, 3-11, 3-12
- OR REPLACE clause
  - of CREATE CONTEXT, 7-261
  - of CREATE DIRECTORY, 7-283
  - of CREATE FUNCTION, 7-286, 7-312
  - of CREATE LIBRARY, 7-316
  - of CREATE OUTLINE, 7-342
  - of CREATE PACKAGE, 7-345
  - of CREATE PACKAGE BODY, 7-349
  - of CREATE PROCEDURE, 7-356
  - of CREATE TRIGGER, 7-428
  - of CREATE TYPE, 7-440
  - of CREATE TYPE BODY, 7-448

- of CREATE VIEW, 7-458
- Oracle reserved words, C-1
- Oracle Tools
  - support of SQL, 1-5
- Oracle8i
  - Enterprise Edition
    - features and functionality, xv
    - features and functionality, xv
    - new features, xvi
- ORDER BY clause
  - of CREATE TABLE, 7-409
  - of queries, 5-21
  - of SELECT, 5-21, 7-571, 7-580
    - with ROWNUM, 2-61
  - of subqueries in CREATE TABLE, 7-409
- ORDER clause
  - of ALTER SEQUENCE. *See* CREATE SEQUENCE
  - of CREATE SEQUENCE, 7-374
- ORDER MEMBER clause
  - of ALTER TYPE, 7-189
  - of CREATE TYPE, 7-443
  - of CREATE TYPE BODY, 7-449
- ORDER methods
  - specifying, 7-189
- ORDERED hint, 2-66
- ORDERED\_PREDICATES hint, 2-65
- ordinal numbers
  - specifying, 2-51
  - spelling out, 2-51
- ORGANIZATION HEAP clause
  - of CREATE TABLE, 7-396
- ORGANIZATION INDEX clause
  - of CREATE TABLE, 7-396
- OUT parameter
  - of CREATE FUNCTION, 7-287
  - of CREATE PROCEDURE, 7-356
- outer joins, 5-23, 7-577
  - restrictions, 5-23
- outlines
  - assigning to a different category, 7-63
  - automatically creating and storing, 7-109
  - creating, 7-342
  - dropping from the database, 7-490
  - enabling and disabling dynamically, 7-342
  - granting

- system privileges on, 7-524
    - rebuilding, 7-63
    - renaming, 7-63
    - replacing, 7-342
    - storing during the session, 7-86
    - storing for the instance, 7-118
    - use by the optimizer, 7-95, 7-118
    - used to generate execution plans, 7-342
- OVERFLOW clause
  - of ALTER INDEX, 7-34
  - of ALTER TABLE, 7-154
  - of CREATE TABLE, 7-397

## P

---

- package bodies
  - creating, 7-348
  - recompiling, 7-64
  - re-creating, 7-349
  - removing from the database, 7-491
- PACKAGE clause
  - of ALTER PACKAGE, 7-65
- packaged procedures
  - dropping, 7-493
- packages
  - associating statistics with, 7-212
  - avoiding run-time compilation, 7-64
  - changing the declaration of, 7-64
  - creating, 7-344
  - disassociating statistics types from, 7-492
  - invoker rights, 7-345
  - recompiling, 7-64
  - redefining, 7-64, 7-345
  - removing from the database, 7-491
  - specifying schema and privileges of, 7-345
  - synonyms for, 7-377
- PACKAGES clause
  - of ASSOCIATE STATISTICS, 7-210, 7-212
- PARALLEL clause
  - of ALTER CLUSTER, 7-3, 7-5
  - of ALTER DATABASE, 7-15
  - of ALTER INDEX, 7-31, 7-36
  - of ALTER MATERIALIZED VIEW, 7-48, 7-51
  - of ALTER MATERIALIZED VIEW LOG, 7-59, 7-60

- of ALTER TABLE, 7-164
- of CREATE CLUSTER, 7-258
- of CREATE INDEX, 7-304
- of CREATE MATERIALIZED VIEW /  
SNAPSHOT, 7-319
- of CREATE MATERIALIZED VIEW LOG /  
SNAPSHOT LOG, 7-334
- of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336
- of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-323
- of CREATE TABLE, 7-388, 7-405
- parallel execution
  - of DDL statements, 7-84
  - of DML statements, 7-84
- PARALLEL hint, 2-67
- PARALLEL\_ADAPTIVE\_MULTI\_USER parameter
  - of ALTER SYSTEM, 7-115
- PARALLEL\_BROADCAST\_ENABLED parameter
  - of ALTER SESSION, 7-92
- PARALLEL\_ENABLE clause
  - of CREATE FUNCTION, 7-288
- PARALLEL\_INDEX hint, 2-67
- PARALLEL\_INSTANCE\_GROUP parameter
  - of ALTER SESSION, 7-92
  - of ALTER SYSTEM, 7-115
- PARALLEL\_MIN\_PERCENT parameter
  - of ALTER SESSION parameter, 7-92
- PARALLEL\_THREADS\_PER\_CPU parameter
  - of ALTER SYSTEM, 7-116
- parameters
  - in syntax diagrams, xxi
  - optional, A-2
  - required, A-1
- PARAMETERS clause
  - of CREATE INDEX, 7-304
- PARTITION ... LOB storage clause
  - of ALTER TABLE, 7-142
- PARTITION BY HASH clause
  - of CREATE TABLE, 7-402
- PARTITION BY RANGE clause
  - of CREATE TABLE, 7-386, 7-403
- PARTITION clause
  - of ANALYZE, 7-203
  - of CREATE INDEX, 7-302
  - of CREATE TABLE, 7-403
  - of DELETE, 7-466
  - of INSERT, 7-542
  - of LOCK TABLE, 7-548
  - of SELECT, 7-575
  - of UPDATE, 7-618
- PARTITION\_VIEW\_ENABLED parameter
  - of ALTER SESSION, 7-93
- partitioned indexes, 2-70, 7-295, 7-302
  - user-defined, 7-302
- partitioned tables, 2-70
- partition-extended table names, 2-70
  - in DML statements, 2-71
  - restrictions on, 2-71
  - syntax, 2-71
- partitioning
  - by range, 7-386
  - clauses
    - of ALTER INDEX, 7-32
    - of ALTER MATERIALIZED VIEW, 7-48,  
7-51
    - of ALTER MATERIALIZED VIEW  
LOG, 7-59, 7-60
    - of ALTER TABLE, 7-155
    - of CREATE MATERIALIZED VIEW /  
SNAPSHOT, 7-319
    - of CREATE MATERIALIZED VIEW LOG /  
SNAPSHOT LOG, 7-334
    - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336
    - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-323
- partitions
  - adding rows to, 7-540
  - allocating extents for, 7-150
  - composite, 2-70
    - specifying, 7-403
  - converting into nonpartitioned tables, 7-162
  - deallocating unused space from, 7-150
  - dropping, 7-160
  - extents
    - allocating new, 7-35
  - hash, 2-70
    - adding, 7-160
    - coalescing, 7-160

- specifying, 7-402
- inserting rows into, 7-542
- LOB storage characteristics of, 7-142
- locking, 7-547
- logging attribute of, 7-395
- logging insert operations, 7-152
- merging, 7-162
- modifying, 7-156
- moving to a different segment, 7-157
- physical attributes
  - changing, 7-145
- range, 2-70
  - adding, 7-158
  - specifying, 7-403
- removing rows from, 7-160, 7-466
- renaming, 7-157
- revising values in, 7-618
- splitting, 7-161
- storage characteristics, 7-394
- tablespace for
  - defining, 7-394
- PASSWORD\_EXPIRE clause
  - of ALTER USER. *See* CREATE USER
  - of CREATE USER, 7-453
- PASSWORD\_GRACE\_TIME parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- PASSWORD\_LIFE\_TIME parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- PASSWORD\_LOCK\_TIME parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- PASSWORD\_REUSE\_MAX parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- PASSWORD\_REUSE\_TIME parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- PASSWORD\_VERIFY\_FUNCTION parameter
  - of ALTER PROFILE, 7-70
  - of CREATE PROFILE, 7-362
- passwords
  - expiration of, 7-453
  - parameters
    - of ALTER PROFILE, 7-363
    - of CREATE PROFILE, 7-360
- PCTFREE parameter
  - of ALTER CLUSTER, 7-4
  - of ALTER INDEX, 7-31, 7-36
  - of ALTER MATERIALIZED VIEW, 7-49
  - of ALTER MATERIALIZED VIEW LOG, 7-58
  - of CREATE CLUSTER, 7-256
  - of CREATE INDEX, 7-300
  - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG. *See* CREATE TABLE.
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT. *See* CREATE TABLE.
  - of CREATE TABLE, 7-393
- PCTINCREASE parameter
  - of STORAGE clause, 7-607
- PCTTHRESHOLD parameter
  - of ALTER TABLE, 7-154
  - of CREATE TABLE, 7-397
- PCTUSED parameter
  - of ALTER CLUSTER, 7-4
  - of ALTER INDEX, 7-31, 7-36
  - of ALTER MATERIALIZED VIEW, 7-49
  - of ALTER MATERIALIZED VIEW LOG, 7-58
  - of CREATE CLUSTER, 7-256
  - of CREATE INDEX. *See* CREATE TABLE
  - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG. *See* CREATE TABLE.
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT. *See* CREATE TABLE.
  - of CREATE TABLE, 7-393
- PCTVERSION parameter
  - of CREATE TABLE, 7-399
  - of LOB storage clause, 7-142
- PERCENT\_RANK function, 4-66
- PERMANENT clause
  - of ALTER TABLESPACE, 7-181
  - of CREATE TABLESPACE, 7-422
- physical attributes clause
  - of a constraint, 7-236
  - of ALTER CLUSTER, 7-2
  - of ALTER INDEX, 7-31, 7-36
  - of ALTER MATERIALIZED VIEW, 7-49
  - of ALTER MATERIALIZED VIEW LOG, 7-58
  - of ALTER TABLE, 7-145



- of CREATE CLUSTER, 7-254
- of CREATE MATERIALIZED VIEW /  
SNAPSHOT, 7-319
- of CREATE MATERIALIZED VIEW LOG /  
SNAPSHOT LOG, 7-334
- of CREATE TABLE, 7-384, 7-393
- plan stability, 7-342
- PLAN\_TABLE sample table, 7-512
- PL/SQL, xv
  - blocks
    - syntax of, xxiii
  - compatibility with earlier releases, 7-93, 7-116
  - program body
    - of CREATE FUNCTION, 7-289
- PLSQL\_V2\_COMPATIBILITY parameter
  - of ALTER SESSION, 7-93
  - of ALTER SYSTEM, 7-116
- PM (P.M.) date format element, 2-46, 2-48
- POWER function, 4-67
- PQ\_DISTRIBUTE hint, 2-67
- PR number format element, 2-41
- PRAGMA clause
  - of ALTER TYPE, 7-189
  - of CREATE TYPE, 7-438, 7-442
- PRAGMA RESTRICT\_REFERENCES, 7-189, 7-442
- precedence
  - of operators, 3-2
- precision
  - number of digits of, 2-4
  - of NUMBER datatype, 2-17
- precompilers
  - Oracle, 1-4
- PRESERVE SNAPSHOT LOG clause
  - of TRUNCATE, 7-613
- PRIMARY KEY clause
  - of constraint\_clause, 7-239
  - of CREATE TABLE, 7-393
- primary key constraints, 7-239
  - enabling, 7-407
  - index on, 7-407
- primary keys
  - generating values for, 7-371
- PRIOR operator, 3-16
- PRIVATE\_SGA parameter
  - of ALTER PROFILE, 7-69
  - of ALTER RESOURCE COST, 7-74
- of CREATE PROFILE, 7-362
- privileges. *See* system privileges or object privileges
- procedures
  - 3GL, calling, 7-316
  - calling, 7-226
  - changing the declaration of, 7-67
  - changing the definition of, 7-67
  - creating, 7-354
  - declaring
    - as a Java method, 7-357
    - as C functions, 7-357
  - executing, 7-226
  - external, 7-354
  - granting
    - system privileges on, 7-524
  - invalidating local objects dependent on, 7-493
  - issuing COMMIT or ROLLBACK  
statements, 7-84
  - naming rules, 2-74
  - privileges executed with, 7-188, 7-441
  - recompiling, 7-67
  - re-creating, 7-356
  - removing from the database, 7-493
  - schema executed in, 7-188, 7-441
  - specification of
    - in CREATE TYPE, 7-442
  - specifying schema and privileges for, 7-357
  - synonyms for, 7-377
- PROFILE clause
  - of ALTER USER. *See* CREATE USER
  - of CREATE USER, 7-453
- profiles
  - assigning to a user, 7-453
  - creating, 7-360
    - examples, 7-363
  - deassigning from users, 7-494
  - granting
    - system privileges on, 7-524
  - modifying, examples, 7-71
  - removing from the database, 7-494
- proxy clause
  - of ALTER USER, 7-194, 7-195
- pseudocolumns, 2-56
  - CURRVAL, 2-56

- LEVEL, 2-58
- NEXTVAL, 2-56
- ROWID, 2-59
- ROWNUM, 2-60
  - uses for, 2-61
- PUBLIC clause
  - of CREATE ROLLBACK SEGMENT, 7-367
  - of CREATE SYNONYM, 7-378
  - of DROP DATABASE LINK, 7-475
- public database links
  - dropping, 7-475
- public rollback segments, 7-367
- public synonyms, 7-378
  - dropping, 7-500
- PURGE SNAPSHOT LOG clause
  - of TRUNCATE, 7-613
- PUSH\_JOIN\_PRED hint, 2-67
- PUSH\_SUBQ hint, 2-67

## Q

---

- Q date format element, 2-46
- queries, 5-19, 7-572
  - comments in, 5-19
  - compound, 5-21
  - correlated
    - left correlation, 7-576
  - defined, 5-19
  - distributed, 5-26
  - grouping returned rows on a value, 7-578
  - hierarchical. *See* hierarchical queries
  - hints in, 5-19
  - join, 5-22
  - locking rows during, 7-580
  - ordering returned rows, 7-580
  - outer joins in, 7-576, 7-577
  - referencing multiple tables, 5-22
  - restricting results of, 7-577
  - select lists of, 5-19
  - selecting from a random sample of rows, 7-574
  - selecting from specified partitions, 7-575
  - sorting results, 5-21
  - syntax, 5-19
  - top-level, 5-19
  - top-N, 2-61

- query rewrite
  - and dimensions, 7-278
  - and function-based indexes, 7-93
  - and rule-based optimization, 7-93
  - consistency level, 7-93, 7-116
  - defined, 7-572
  - disabling, 7-93, 7-116
  - enabling, 7-93, 7-116
- QUERY REWRITE system privilege, 7-523, 7-526
- QUERY\_REWRITE\_ENABLED parameter
  - of ALTER SESSION, 7-93
  - of ALTER SYSTEM, 7-116
- QUERY\_REWRITE\_INTEGRITY parameter
  - of ALTER SESSION, 7-93
  - of ALTER SYSTEM, 7-116
- QUOTA clause
  - of ALTER USER. *See* CREATE USER
  - of CREATE USER, 7-453

## R

---

- range conditions, 5-18
- range partition
  - adding, 7-158
  - creating, 7-403
- RANK function, 4-67
- RATIO\_TO\_REPORT function, 4-68
- RAW data
  - converting from CHAR data, 2-23
- RAW datatype, 2-22
- RAWTOHEX function, 2-37, 4-69
- READ object privilege, 7-535
- READ ONLY clause
  - of ALTER TABLESPACE, 7-181
  - of SET TRANSACTION, 7-602
- READ WRITE clause
  - of ALTER TABLESPACE, 7-181
  - of SET TRANSACTION, 7-603
- REAL datatype
  - ANSI, 2-28
- REBUILD clause
  - of ALTER INDEX, 7-32, 7-37
  - of ALTER OUTLINE, 7-63
- REBUILD COMPRESS clause
  - of ALTER INDEX, 7-39

- REBUILD COMPUTE STATISTICS clause
  - of ALTER INDEX, 7-39
- REBUILD LOGGING clause
  - of ALTER INDEX, 7-39
- REBUILD NOCOMPRESS clause
  - of ALTER INDEX, 7-39
- REBUILD NOLOGGING clause
  - of ALTER INDEX, 7-39
- REBUILD NOREVERSE clause
  - of ALTER INDEX, 7-38
- REBUILD ONLINE clause
  - of ALTER INDEX, 7-39
- REBUILD PARAMETERS clause
  - of ALTER INDEX, 7-40
- REBUILD PARTITION clause
  - of ALTER INDEX, 7-38
- REBUILD REVERSE clause
  - of ALTER INDEX, 7-38
- REBUILD SUBPARTITION clause
  - of ALTER INDEX, 7-38
- REBUILD TABLESPACE clause
  - of ALTER INDEX, 7-38
- REBUILD UNUSABLE LOCAL INDEXES clause
  - of ALTER TABLE, 7-157
- RECOVER AUTOMATIC clause
  - of ALTER DATABASE, 7-13
- RECOVER CANCEL clause
  - of ALTER DATABASE, 7-7, 7-14
- RECOVER clause
  - of ALTER DATABASE, 7-7, 7-12
- RECOVER CONTINUE clause
  - of ALTER DATABASE, 7-7, 7-14
- RECOVER DATABASE clause
  - of ALTER DATABASE, 7-7, 7-13
- RECOVER DATAFILE clause
  - of ALTER DATABASE, 7-7, 7-14
- RECOVER LOGFILE clause
  - of ALTER DATABASE, 7-7, 7-14
- RECOVER MANAGED STANDBY DATABASE clause
  - of ALTER DATABASE, 7-8
- RECOVER STANDBY DATABASE clause
  - of ALTER DATABASE, 7-13
- RECOVER STANDBY DATAFILE clause
  - of ALTER DATABASE, 7-14
- RECOVER STANDBY TABLESPACE clause
  - of ALTER DATABASE, 7-14
- RECOVER TABLESPACE clause
  - of ALTER DATABASE, 7-7, 7-14
- RECOVERABLE, 7-37, 7-396
  - See also* LOGGING clause
- recovery
  - distributed, enabling, 7-106
  - of database, 7-7
- RECOVERY\_CATALOG\_OWNER role, 7-530
- redo logs
  - adding, 7-18
  - automatic archiving of, 7-103
  - automatic name generation, 7-13
  - disabling specified threads in a parallel server, 7-16
  - dropping, 7-19
  - enabling specified threads in a parallel server, 7-15
  - members
    - adding to existing groups, 7-18
    - dropping, 7-19
    - renaming, 7-15
  - reusing, 7-517
  - size of, 7-517
  - specifying, 7-516
    - for media recovery, 7-14
  - switching groups, 7-107
- REF columns
  - specifying, 7-393
  - specifying from table or column level, 7-393
- REF function, 4-70
- REFERENCES clause
  - of constraint\_clause, 7-241
  - of CREATE TABLE, 7-393
- REFERENCES object privilege, 7-535
- references to objects. *See* REFS
- REFERENCING clause
  - of CREATE TRIGGER, 7-427, 7-432
- referential integrity constraints, 7-239, 7-241
- REFRESH clause
  - of ALTER MATERIALIZED VIEW, 7-48, 7-52
  - of CREATE MATERIALIZED VIEW / SNAPSHOT, 7-319
- REFRESH COMPLETE clause

- of ALTER MATERIALIZED VIEW, 7-52
- of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-324
- REFRESH FAST clause
  - of ALTER MATERIALIZED VIEW, 7-52
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-324
- REFRESH FORCE clause
  - of ALTER MATERIALIZED VIEW, 7-53
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-324
- REFRESH ON COMMIT clause
  - of ALTER MATERIALIZED VIEW, 7-53
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-324
- REFRESH ON DEMAND clause
  - of ALTER MATERIALIZED VIEW, 7-53
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-324
- REFs, 2-31, 7-242
  - as containers for OIDs, 2-31
  - dangling, 7-206
  - validating, 7-206
- REFTOHEX function, 4-70
- REGR\_AVGX function, 4-71
- REGR\_AVGY function, 4-71
- REGR\_COUNT function, 4-71
- REGR\_INTERCEPT function, 4-71
- REGR\_R2 function, 4-71
- REGR\_SLOPE function, 4-71
- REGR\_SXX function, 4-71
- REGR\_SXY function, 4-71
- REGR\_SYY function, 4-71
- relational tables
  - creating, 7-381
- RELY clause
  - of constraint\_clause, 7-244
- REMOTE\_DEPENDENCIES\_MODE parameter
  - of ALTER SESSION, 7-94
  - of ALTER SYSTEM, 7-117
- REMOTE\_LOGIN\_PASSWORDFILE parameter
  - and control files, 7-264
  - and databases, 7-269
- RENAME clause
  - of ALTER INDEX, 7-40
  - of ALTER OUTLINE, 7-63
  - of ALTER TABLE, 7-153
- RENAME DATAFILE clause
  - of ALTER TABLESPACE, 7-179
- RENAME FILE clause
  - of ALTER DATABASE, 7-6, 7-15
- RENAME GLOBAL\_NAME clause
  - of ALTER DATABASE, 7-15
- RENAME PARTITION clause
  - of ALTER INDEX, 7-33, 7-41
  - of ALTER TABLE, 7-157
- RENAME statement, 7-554
- RENAME SUBPARTITION clause
  - of ALTER INDEX, 7-33, 7-41
  - of ALTER TABLE, 7-157
- REPLACE AS OBJECT clause
  - of ALTER TYPE, 7-188
- REPLACE function, 4-78
- reserved words, 2-73, C-1
- RESET COMPATIBILITY clause
  - of ALTER DATABASE, 7-15
- RESETLOGS parameter
  - of CREATE CONTROLFILE, 7-265
- RESOLVE clause
  - of ALTER JAVA CLASS, 7-46
  - of CREATE JAVA, 7-312
- RESOLVER clause
  - of ALTER JAVA CLASS, 7-46
  - of ALTER JAVA SOURCE, 7-46
  - of CREATE JAVA, 7-314
- resource parameters
  - of CREATE PROFILE, 7-359
- RESOURCE role, 7-530
- RESOURCE\_LIMIT parameter
  - of ALTER SYSTEM, 7-117
- RESOURCE\_MANAGER\_PLAN parameter
  - of ALTER SYSTEM, 7-117
- RESTRICT\_REFERENCES pragma
  - of ALTER TYPE, 7-189
- restricted rowids, 2-26
  - compatibility and migration of, 2-27
- RESTRICTED SESSION system privilege, 7-525
- RESUME clause
  - of ALTER SYSTEM, 7-108
- RETURN clause

- of CREATE FUNCTION, 7-288
  - of CREATE OPERATOR, 7-340
  - of CREATE TYPE BODY, 7-449
- RETURNING clause
  - of DELETE, 7-467
  - of INSERT, 7-540, 7-544
  - of UPDATE, 7-616, 7-620
- REUSE clause
  - of CREATE CONTROLFILE, 7-264
  - of filespec clause, 7-517
- REUSE STORAGE clause
  - of TRUNCATE, 7-613
- REVERSE clause
  - of CREATE INDEX, 7-300
- reverse indexes, 7-300
- REVOKE CONNECT THROUGH clause
  - of ALTER USER, 7-194, 7-195
- REVOKE schema\_object\_privileges
  - statement, 7-559
- REVOKE system\_privileges\_and\_roles
  - statement, 7-556
- REWRITE hint, 2-65
- RM date format element, 2-46
- RN number format element, 2-41
- RNDS parameter
  - of PRAGMA RESTRICT\_REFERENCES, 7-189
- RNPS parameter
  - of PRAGMA RESTRICT\_REFERENCES, 7-189
- roles
  - authorization
    - by a password, 7-366
    - by an external service, 7-366
    - by the database, 7-366
    - by the enterprise directory service, 7-366
  - changing, 7-76
  - creating, 7-365
  - disabling
    - for the current session, 7-600, 7-601
  - enabling
    - for the current session, 7-600, 7-601
  - granting, 7-519
    - system privileges on, 7-525
    - to a user, 7-520
    - to all users, 7-520
    - to another role, 7-520
    - to PUBLIC, 7-520
  - removing from the database, 7-495
  - revoking, 7-556
    - from another role, 7-495, 7-557
    - from PUBLIC, 7-557
    - from users, 7-495, 7-557
- rollback segments
  - bringing online, 7-78, 7-368
  - changing storage characteristics, 7-78
  - creating, 7-367
  - granting
    - system privileges on, 7-525
  - multiple, 7-368
  - public, 7-367
  - reducing size, 7-78
  - removing from the database, 7-496
  - specifying optimal size of, 7-608
  - specifying tablespaces for, 7-367
  - SQL examples, 7-368
  - storage characteristics, 7-368, 7-605
  - taking offline, 7-78
- ROLLBACK statement, 7-564
- ROLLUP operation
  - of queries and subqueries, 7-578
- ROUND function
  - date function, 4-79
  - format models, 4-117
  - number function, 4-79
- routines
  - calling, 7-226
  - executing, 7-226
- ROW EXCLUSIVE lock mode, 7-548
- ROW SHARE lock mode, 7-548
- ROW\_NUMBER function, 4-80
- ROWID datatype, 2-26
- ROWID hint, 2-65
- ROWID pseudocolumn, 2-26, 2-27, 2-59
- rowids
  - block portion of, 2-26
  - description of, 2-26
  - extended, 2-27
    - not directly available, 2-27
  - file portion of, 2-26
  - nonphysical, 2-27
  - of foreign tables, 2-27

- of index-organized tables, 2-27
- restricted, 2-26
  - compatibility and migration of, 2-27
- row portion of, 2-26
- uses for, 2-60
- ROWIDTOCHAR function, 2-37, 4-81
- ROWNUM pseudocolumn, 2-60
  - uses for, 2-61
- rows
  - adding to a table, 7-540
  - allowing movement of between partitions, 7-383
  - inserting
    - into partitions, 7-542
    - into remote databases, 7-542
    - into subpartitions, 7-542
  - movement between partitions, 7-404
  - order of storage, 7-396
  - removing
    - from a cluster, 7-611
    - from a table, 7-611
    - from partitions and subpartitions, 7-466
    - from tables and views, 7-465
  - selecting in hierarchical order, 5-20
  - specifying constraints on, 7-241
  - stored in ascending order, 7-244
  - storing if in violation of constraints, 7-163, 7-245, 7-407
- RPAD function, 4-82
- RR date format element, 2-46, 2-49
  - interpreting, 2-49
- RRRR date format element, 2-46
- RTRIM function, 4-82
- RULE hint, 2-64
- run-time compilation
  - avoiding, 7-67, 7-198

## S

---

- S number format element, 2-41
- SALES sample table, 4-3
- SAMPLE clause
  - of SELECT, 7-574
  - of SELECT and subqueries, 7-570
- SAVEPOINT statement, 7-567

- savepoints
  - erasing, 7-230
  - rolling back to, 7-565
  - specifying, 7-567
- scale
  - greater than precision, 2-18
  - negative, 2-18
  - of NUMBER datatype, 2-17
- SCC date format element, 2-46
- SCHEMA clause
  - of CREATE JAVA, 7-313
- schema objects, 2-68
  - auditing
    - by access, 7-222
    - by session, 7-222
    - options, 7-223
    - successful
      - SQL statements on, 7-222
  - defining default buffer pool for, 7-609
  - dropping, 7-509
  - in other schemas, 2-79
  - list of, 2-68
  - name resolution, 2-78
  - namespaces, 2-73
  - naming
    - examples, 2-75
    - guidelines, 2-76
    - rules, 2-72
  - object types, 2-30
  - on remote databases, 2-79
  - partitioned indexes, 2-70
  - partitioned tables, 2-70
  - parts of, 2-70
  - reauthorizing, 6-2
  - recompiling, 6-2
  - referring to, 2-76, 7-87
  - remote, accessing, 7-273
  - stopping auditing of, 7-552
- schemas
  - changing for a session, 7-87
  - creating, 7-369
  - definition of, 2-68
- scientific notation, 2-42
- SCOPE clause
  - of column ref constraints, 7-242

- scope constraints, 7-242
- segment attributes clause
  - of CREATE TABLE, 7-383, 7-393
- SELECT
  - object privilege, 7-535
  - statement, 7-569
- SELECT ANY SEQUENCE system privilege, 7-525
- SELECT ANY TABLE system privilege, 7-527
- select lists, 5-19
  - ordering, 5-21
- SELECT statement, 5-19
- SELECT\_CATALOG\_ROLE role, 7-530
- self joins, 5-23
- sequences, 2-56, 7-372
  - accessing values of, 7-372
  - changing
    - the increment value, 7-81
    - the number of cached values, 7-81
  - creating, 7-371
  - creating without limit, 7-373
  - granting
    - system privileges on, 7-525
  - how to use, 2-57
  - incrementing, 7-371, 7-373
  - maximum value
    - setting or changing, 7-81
  - minimum value
    - setting or changing, 7-81
  - ordering values, 7-81
  - recycling values, 7-81
  - removing from the database, 7-497
  - renaming, 7-554
  - restarting, 7-497
    - at a different number, 7-82
    - at a predefined limit, 7-373
  - reusing, 7-372
  - stopping at a predefined limit, 7-373
  - synonyms for, 7-377
  - where to use, 2-57
- SERVERERROR event
  - triggers on, 7-431
- service name
  - of remote database, 7-275
- session control statements, 6-5
  - PL/SQL support of, 6-5
- session locks
  - releasing, 7-107
- SESSION\_CACHED\_CURSORS parameter
  - of ALTER SESSION, 7-94
- SESSION\_ROLES view, 7-600
- sessions
  - calculating resource cost limits, 7-73
  - changing resource cost limits, 7-73
  - disconnecting, 7-106
  - global name resolution for, 7-88
  - granting
    - system privileges on, 7-525
  - limiting resource costs, 7-73
  - modifying characteristics of, 7-86
  - number of concurrent, 7-111
  - object cache, 7-91
  - restricted, 7-106
  - terminating, 7-107
- SESSIONS\_PER\_USER parameter
  - of ALTER PROFILE, 7-69
  - of CREATE PROFILE, 7-361
- SET clause
  - of ALTER SESSION, 7-86
  - of ALTER SYSTEM, 7-108
  - of UPDATE, 7-619
- SET CONSTRAINT(S) statement, 7-598
- SET DATABASE clause
  - of CREATE CONTROLFILE, 7-264
- set operators, 3-12, 7-579
- SET ROLE statement, 7-600
- SET STATEMENT\_ID clause
  - of EXPLAIN PLAN, 7-513
- SET TRANSACTION statement, 7-602
- SET UNUSED clause
  - of ALTER TABLE, 7-147
- SGA. *See* system global area
- SHARE ROW EXCLUSIVE lock mode, 7-548
- SHARE UPDATE lock mode, 7-548
- SHARED clause
  - of CREATE DATABASE LINK, 7-274
- shared server processes
  - creating additional, 7-114
  - terminating, 7-114
- SHRINK clause
  - of ALTER ROLLBACK SEGMENT, 7-79

- SHUTDOWN event
  - triggers on, 7-431
- SIGN function, 4-83
- simple comparison conditions, 5-16
- simple expressions, 5-3
- SIN function, 4-83
- SINGLE TABLE clause
  - of CREATE CLUSTER, 7-257
- single-row functions, 4-4
  - miscellaneous, 4-6
- SINH function, 4-84
- SIZE clause
  - of ALTER CLUSTER, 7-4
  - of CREATE CLUSTER, 7-256
  - of filespec clause, 7-517
- SKIP\_UNUSABLE\_INDEXES parameter
  - of ALTER SESSION, 7-94
- SMALLINT datatype
  - ANSI, 2-28
  - DB2, 2-29
  - SQL/DS, 2-29
- snapshot logs. *See* materialized view logs
- snapshots. *See* materialized views
- SNMPAGENT role, 7-530
- SOME operator, 3-6
- SORT\_AREA\_RETAINED\_SIZE parameter
  - of ALTER SESSION, 7-95
  - of ALTER SYSTEM, 7-118
- SORT\_AREA\_SIZE parameter
  - of ALTER SESSION, 7-95
  - of ALTER SYSTEM, 7-118
- SORT\_MULTIBLOCK\_READ\_COUNT parameter
  - of ALTER SESSION, 7-95
  - of ALTER SYSTEM, 7-118
- SOUNDEX function, 4-84
- SP date format element suffix, 2-51
- SPECIFICATION clause
  - of ALTER PACKAGE, 7-65
- spelled numbers
  - specifying, 2-51
- SPLIT PARTITION clause
  - of ALTER INDEX, 7-34, 7-41
  - of ALTER TABLE, 7-161
- SPTH date format element suffix, 2-51
- SQL
  - description of, 1-2
  - embedded, 1-4
  - functions, 4-2
  - keywords, A-1
  - Oracle Tools support of, 1-5
  - parameters, A-1
  - standards, 1-2, B-1
  - statements
    - auditing, 7-215
    - determining the cost of, 7-512
    - syntax, 7-1, A-1
  - SQL function
    - ACOS, 4-13
  - SQL functions
    - ABS, 4-12
    - ADD\_MONTHS, 4-13
    - aggregate, 4-6
    - analytic, 4-7
    - ASCII, 4-14
    - ASIN, 4-15
    - ATAN, 4-15
    - ATAN2, 4-16
    - AVG, 4-16
    - BFILENAME, 4-17
    - CEIL, 4-18
    - character
      - returning character values, 4-4
      - returning number values, 4-5
    - CHARTOROWID, 4-19
    - CHR, 4-19
    - CONCAT, 4-20
    - conversion, 4-5
    - CONVERT, 4-21
    - CORR, 4-22
    - COS, 4-23
    - COSH, 4-24
    - COUNT, 4-24
    - CUME\_DIST, 4-28
    - date, 4-5
    - DENSE\_RANK, 4-29
    - DEREF, 4-30
    - DUMP, 4-31
    - EMPTY\_BLOB, 4-33
    - EMPTY\_CLOB, 4-33
    - EXP, 4-33



FIRST\_VALUE, 4-34  
 FLOOR, 4-35  
 GREATEST, 4-36  
 GROUPING, 4-36  
 HEXTORAW, 4-37  
 INITCAP, 4-38  
 INSTR, 4-38  
 INSTRB, 4-39  
 LAG, 4-40  
 LAST\_DAY, 4-41  
 LAST\_VALUE, 4-42  
 LEAD, 4-44  
 LEAST, 4-45  
 LENGTH, 4-45  
 LENGTHB, 4-46  
 linear regression, 4-71  
 LN, 4-46  
 LOG, 4-47  
 LOWER, 4-47  
 LPAD, 4-48  
 LTRIM, 4-48  
 MAKE\_REF, 4-49  
 MAX, 4-50  
 MIN, 4-52  
 MOD, 4-53  
 MONTHS\_BETWEEN, 4-54  
 NEW\_TIME, 4-55  
 NEXT\_DAY, 4-55  
 NLS\_CHARSET\_DECL\_LEN, 4-56  
 NLS\_CHARSET\_ID, 4-56  
 NLS\_CHARSET\_NAME, 4-57  
 NLS\_INITCAP, 4-58  
 NLS\_LOWER, 4-59  
 NLS\_UPPER, 4-60  
 NLSSORT, 4-59  
 NLV2, 4-65  
 NTILE, 4-61  
 number, 4-4  
 NUMTODSINTERVAL, 4-62  
 NUMTOYMINTERVAL, 4-63  
 NVL, 4-64  
 object reference, 4-12  
 PERCENT\_RANK, 4-66  
 POWER, 4-67  
 RANK, 4-67  
 RATIO\_TO\_REPORT, 4-68  
 RAWTOHEX, 4-69  
 REF, 4-70  
 REFTOHEX, 4-70  
 REGR\_AVGX, 4-71  
 REGR\_AVGY, 4-71  
 REGR\_COUNT, 4-71  
 REGR\_INTERCEPT, 4-71  
 REGR\_R2, 4-71  
 REGR\_SLOPE, 4-71  
 REGR\_SXX, 4-71  
 REGR\_SXY, 4-71  
 REGR\_SYY, 4-71  
 REPLACE, 4-78  
 ROUND (date), 4-79  
 ROUND (number), 4-79  
 ROW\_NUMBER, 4-80  
 ROWIDTOCHAR, 4-81  
 RPAD, 4-82  
 RTRIM, 4-82  
 SIGN, 4-83  
 SIN, 4-83  
 single-row, 4-4  
     miscellaneous, 4-6  
 SINH, 4-84  
 SOUNDEX, 4-84  
 SQRT, 4-85  
 STDDEV, 4-86  
 STDDEVP, 4-87  
 STDDEVS, 4-88  
 SUBSTR, 4-90  
 SUBSTRB, 4-91  
 SUM, 4-91  
 SYS\_CONTEXT, 4-93  
 SYS\_GUID, 4-97  
 SYSDATE, 4-98  
 TAN, 4-98  
 TANH, 4-99  
 TO\_CHAR (date), 4-99  
 TO\_CHAR (number), 4-100  
 TO\_DATE, 4-101  
 TO\_LOB, 4-102  
 TO\_MULTI\_BYTE, 4-103  
 TO\_NUMBER, 4-103  
 TO\_SINGLE\_BYTE, 4-104

- TRANSLATE, 4-105
- TRANSLATE...USING, 4-106
- TRIM, 4-107
- TRUNC (date), 4-109
- TRUNC (number), 4-108
- UID, 4-109
- UPPER, 4-110
- USER, 4-110
- USERENV, 4-111
- VALUE, 4-112
- VARIANCE, 4-115
- VARP, 4-113
- VARS, 4-114
- VSIZE, 4-116
- SQL statements
  - auditing
    - by access, 7-215
    - by proxy, 7-214
    - by session, 7-214
    - by user, 7-214
    - stopping, 7-550
    - successful, 7-215
  - determining the execution plan for, 7-512
  - rolling back, 7-564
  - tracking the occurrence in a session, 7-213
  - undoing, 7-564
- SQL\_TRACE parameter
  - of ALTER SESSION, 7-95
- SQL92, 1-2
  - Oracle compliance with, B - 2
- SQL/DS datatypes, 2-28
  - conversion to Oracle datatypes, 2-29
  - restrictions on, 2-30
- SQRT function, 4-85
- SS date format element, 2-46
- SSSS date format element, 2-46
- standalone procedures
  - dropping, 7-493
- standard SQL, B - 1
  - Oracle extensions to, B - 5
- standby control file
  - creating, 7-20
- standby database
  - activating, 7-11
  - designing media recovery, 7-12
  - mounting, 7-11
  - recovering, 7-13, 7-14
- STANDBY\_ARCHIVE\_DEST parameter
  - of ALTER SYSTEM, 7-118
- STAR hint, 2-66
- STAR\_TRANSFORMATION hint, 2-67
- STAR\_TRANSFORMATION\_ENABLED parameter
  - of ALTER SESSION, 7-95
- START WITH clause
  - of ALTER MATERIALIZED VIEW...REFRESH, 7-53
  - of CREATE SEQUENCE, 7-373
  - of SELECT, 7-577
  - of SELECT and subqueries, 7-571
- STARTUP event
  - triggers on, 7-431
- STATIC clause
  - of ALTER TYPE, 7-188
  - of CREATE TYPE, 7-442
  - of CREATE TYPE BODY, 7-448
- statistics
  - computing exactly, 7-204
  - deleting from the data dictionary, 7-206
  - estimating, 7-204
  - forcing disassociation, 7-471
  - on indexes, 7-301
  - user-defined
    - dropping, 7-478, 7-480, 7-482, 7-492, 7-501, 7-506
- statistics types
  - associating
    - with columns, 7-211
    - with datatypes, 7-212
    - with domain indexes, 7-212
    - with functions, 7-212
    - with indextypes, 7-212
    - with packages, 7-212
  - disassociating
    - from columns, 7-470
    - from domain indexes, 7-470
    - from functions, 7-470
    - from indextypes, 7-470
    - from packages, 7-470
    - from types, 7-470
- STDDEV function, 4-86

STDDEVP function, 4-87  
 STDDEVS function, 4-88  
 storage characteristics  
   resetting, 7-611  
 STORAGE clause, 7-605  
   of ALTER CLUSTER, 7-4  
   of ALTER INDEX, 7-31, 7-36  
   of ALTER MATERIALIZED VIEW, 7-49  
   of ALTER MATERIALIZED VIEW LOG, 7-58  
   of ALTER ROLLBACK SEGMENT, 7-78, 7-79  
   of CREATE CLUSTER, 7-256  
   of CREATE INDEX, 7-300  
   of CREATE MATERIALIZED VIEW LOG /  
     SNAPSHOT LOG, 7-334  
   of CREATE MATERIALIZED VIEW LOG /  
     SNAPSHOT LOG. *See* CREATE TABLE  
   of CREATE MATERIALIZED VIEW /  
     SNAPSHOT. *See* CREATE TABLE.  
   of CREATE ROLLBACK SEGMENTS, 7-368  
   of CREATE TABLE, 7-384, 7-394  
   of CREATE TABLESPACE, 7-420  
 STORAGE IN ROW clause  
   of ALTER TABLE, 7-141  
 STORE IN DEFAULT clause  
   of CREATE INDEX, 7-303  
 STORE IN tablespace clause  
   of CREATE INDEX, 7-303  
 stored functions, 7-285  
 Structured Query Language. *See* SQL  
 SUBPARTITION BY HASH clause  
   of CREATE TABLE, 7-387, 7-403  
 SUBPARTITION clause  
   of ANALYZE, 7-203  
   of CREATE INDEX, 7-303  
   of CREATE TABLE, 7-404  
   of DELETE, 7-466  
   of INSERT, 7-542  
   of LOCK TABLE, 7-548  
   of SELECT, 7-575  
   of UPDATE, 7-618  
 subpartition-extended table names, 2-70  
   in DML statements, 2-71  
   restrictions on, 2-71  
   syntax, 2-71  
 subpartitions  
   adding, 7-156  
   adding rows to, 7-540  
   allocating extents for, 7-150, 7-157  
   coalescing, 7-156  
   converting into nonpartitioned tables, 7-162  
   creating, 7-387, 7-404  
   deallocating unused space from, 7-150, 7-157  
   inserting rows into, 7-542  
   locking, 7-547  
   logging insert operations, 7-152  
   moving to a different segment, 7-158  
   physical attributes  
     changing, 7-145  
   removing rows from, 7-160, 7-466  
   renaming, 7-157  
   revising values in, 7-618  
   specifying, 7-403  
 SUBPARTITIONS clause  
   of CREATE TABLE, 7-403  
 subqueries, 5-19, 7-569, 7-572  
   containing subqueries, 5-25  
   correlated, 5-25  
   defined, 5-19  
   to insert table data, 7-409  
 SUBSTR function, 4-90  
 SUBSTRB function, 4-91  
 subtotal values  
   deriving, 7-578  
 SUM function, 4-91  
 SUSPEND clause  
   of ALTER SYSTEM, 7-107  
 sustained standby recovery mode, 7-14  
   terminating, 7-15  
   timeout period, 7-14  
 SWITCH LOGFILE clause  
   of ALTER SYSTEM, 7-107  
 SYEAR date format element, 2-46  
 synonyms  
   changing the definition of, 7-500  
   creating, 7-377  
   granting  
     system privileges on, 7-526  
   local, 7-379  
   private, dropping, 7-500  
   public, 7-378

- dropping, 7-500
- remote, 7-379
- removing from the database, 7-500
- renaming, 7-554
- synonyms for, 7-377
- syntax diagrams, A-1
  - explanation of, xx
  - keywords, xxi
  - loops, A-2
  - multipart diagrams, A-3
  - parameters, xxi
- SYS schema
  - database triggers stored in, 7-433
  - functions stored in, 7-433
- SYS\_CONTEXT function, 4-93
- SYS\_GUID function, 4-97
- SYSDATE function, 4-98
- SYSDBA system privilege, 7-529
- SYSOPER system privilege, 7-529
- system control statements, 6-5
  - PL/SQL support of, 6-5
- system date
  - altering, 7-110
- system events
  - attributes of, 7-433
  - triggers on, 7-431
- system global area
  - flushing, 7-107
  - updating, 7-105
- system privileges
  - granting, 7-365, 7-519
    - to a role, 7-520
    - to a user, 7-520
    - to all users, 7-520
    - to PUBLIC, 7-520
  - list of, 7-521
  - revoking, 7-556
    - from a role, 7-557
    - from a user, 7-556
    - from PUBLIC, 7-557
    - from users, 7-557
- YYYY date format element, 2-46

## T

---

- TABLE clause
  - of DELETE, 7-467
  - of INSERT, 7-542
  - of SELECT, 7-576
  - of TRUNCATE, 7-612
  - of UPDATE, 7-617, 7-618
- table constraints
  - defined, 7-233
  - of ALTER TABLE, 7-141
  - of CREATE TABLE, 7-393
- table locks
  - disabling, 7-165
  - duration of, 7-547
  - enabling, 7-165
  - EXCLUSIVE, 7-548
  - modes of, 7-548
  - on partitions, 7-548
  - on remote database, 7-548
  - on subpartitions, 7-548
  - and queries, 7-547
  - ROW EXCLUSIVE, 7-548
  - ROW SHARE, 7-548
  - SHARE, 7-548
  - SHARE ROW EXCLUSIVE, 7-548
  - SHARE UPDATE, 7-548
- table REF constraints, 7-234, 7-242
  - of ALTER TABLE, 7-140
  - of CREATE TABLE, 7-393
- tables
  - adding rows to, 7-540
  - aliases, 2-82
    - in CREATE INDEX, 7-298
    - in DELETE, 7-467
  - allocating extents for, 7-150
  - assigning to a cluster, 7-401
  - changing degree of parallelism on, 7-164
  - changing existing values in, 7-616
  - collecting modification statistics on, 7-152
  - collecting statistics on, 7-203
  - creating, 7-390
    - multiple, 7-369
  - creating comments about, 7-228
  - deallocating unused space from, 7-150

- default physical attributes
  - changing, 7-145
- degree of parallelism
  - specifying, 7-390
- disassociating statistics types from, 7-501
- dropping
  - along with cluster, 7-473
  - along with owner, 7-510
  - indexes of, 7-501
  - partitions of, 7-501
- granting
  - system privileges on, 7-526
- index-organized
  - creating, 7-396
  - overflow segment for, 7-397
  - space in index block, 7-397
- inserting rows with a subquery, 7-409
- LOB storage of, 7-394
- locking, 7-547
- logging
  - insert operations, 7-152
  - table creation, 7-395
- migrated and chained rows in, 7-207
- moving to a new segment, 7-144
- nested
  - creating, 7-444
  - storage characteristics, 7-400
- object
  - creating, 7-381
- order of row storage, 7-396
- ordering rows from, 7-409
- parallel creation of, 7-405
- parallelism
  - setting default degree, 7-405
- partition attributes of, 7-155
- partitioned
  - allowing rows to move between partitions, 7-164
  - default attributes of, 7-155
- partitioning of, 2-70, 7-390
- physical attributes
  - changing, 7-145
  - specifying, 7-393
- relational
  - creating, 7-381
- remote, accessing, 7-273
- removing from the database, 7-501
- removing rows from, 7-465
- renaming, 7-153, 7-554
- restricting
  - records per block, 7-153
  - references to, 7-242
- retrieving data from, 7-572
- rows
  - ordering by primary key, 7-396
- saving blocks in a cache, 7-151, 7-408
- SQL examples, 7-410
- storage characteristics, 7-605
  - defining, 7-390, 7-394
- subpartition attributes of, 7-155
- synonyms for, 7-377
- tablespace for
  - defining, 7-390, 7-394
- temporary
  - duration of data, 7-401
  - session-specific, 7-391
  - transaction specific, 7-391
- unclustering, 7-472
- validating structure of, 7-206
- with unusable indexes, 7-94
- TABLESPACE clause
  - of CREATE CLUSTER, 7-257
  - of CREATE INDEX, 7-300
  - of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-322
  - of CREATE ROLLBACK SEGMENTS, 7-367
  - of CREATE TABLE, 7-394
- tablespaces, 7-180
  - allocating space for users, 7-453
  - allowing write operations on, 7-181
  - backing up datafiles of, 7-180
  - bringing online, 7-180, 7-422
  - coalescing free extents, 7-181
  - converting
    - from permanent to temporary, 7-181
    - from temporary to permanent, 7-181
  - creating, 7-420
  - datafile

- adding, 7-179
  - renaming, 7-179
- default storage characteristics, 7-605
- defining as read only, 7-181
- designing media recovery, 7-12
- dropping the contents of, 7-504
- enable autoextension of, 7-179
- extent management of, 7-425
- extent size of, 7-421
- granting
  - system privileges on, 7-527
- locally managed, 7-422, 7-606
  - temporary, 7-425
- logging attribute of, 7-178, 7-421
- managed using dictionary tables, 7-422
- managing extents of, 7-422
- of session duration, 7-424
- permanent objects in, 7-422
- recovering, 7-14
- removing from the database, 7-503
- size of free extents in, 7-180
- specifying
  - datafiles for, 7-421
  - for a user, 7-453
  - for index rebuild, 7-145
- taking offline, 7-180, 7-422
- tempfile
  - adding, 7-179
- temporary
  - creating, 7-424
  - specifying for a user, 7-453
- temporary objects in, 7-422

TAN function, 4-98

TANH function, 4-99

TEMPFILE clause
 

- of ALTER DATABASE, 7-8, 7-17
- of CREATE TEMPORARY TABLESPACE, 7-425

tempfiles
 

- automatic extension of, 7-425
- bringing online, 7-17
- disabling automatic extension, 7-17
- dropping, 7-17
- enabling automatic extension, 7-17
- modifying, 7-16
- resizing, 7-17
- reusing, 7-517
- size of, 7-517
- specifying, 7-425, 7-516
- taking offline, 7-17

TEMPORARY clause
 

- of ALTER TABLESPACE, 7-181
- of CREATE TABLESPACE, 7-422

temporary tables
 

- creating, 7-390, 7-391
- session-specific, 7-391
- transaction-specific, 7-391

TEMPORARY TABLESPACE clause
 

- of ALTER USER. *See* CREATE USER
- of CREATE USER, 7-453

temporary tablespaces
 

- creating, 7-424
- specifying for a user, 7-453
- SQL examples, 7-425

text
 

- conventions, xx
- date and number formats, 2-38
- in SQL syntax, 2-2
- properties of CHAR and VARCHAR2 datatypes, 2-2
- syntax of, 2-2

text date format element, 2-46

TH date format element suffix, 2-51

THSP date format element suffix, 2-51

TIME datatype (SQL/DS or DB2), 2-30

TIMED\_OS\_STATISTICS parameter
 

- of ALTER SYSTEM, 7-118

TIMED\_STATISTICS parameter
 

- of ALTER SESSION, 7-95
- of ALTER SYSTEM, 7-118

TIMESTAMP datatype (SQL/DS or DB2), 2-30

TM number format element, 2-41

TO PUBLIC clause
 

- of GRANT object\_privileges, 7-534
- of GRANT system\_privileges\_and\_roles, 7-520

TO role clause
 

- of GRANT object\_privileges, 7-534
- of GRANT system\_privileges\_and\_roles, 7-520

TO SAVEPOINT clause
 

- of ROLLBACK, 7-565

TO user clause

- of GRANT object\_privileges, 7-534
- of GRANT system\_privileges\_and\_roles, 7-520
- TO\_CHAR
  - date conversion function, 4-99
  - number conversion function, 4-100
- TO\_CHAR function, 2-37, 2-40, 2-45, 2-51
- TO\_DATE function, 2-37, 2-45, 2-49, 2-51, 4-101
- TO\_LOB function, 2-37, 4-102
- TO\_MULTI\_BYTE function, 4-103
- TO\_NUMBER function, 2-37, 2-40, 4-103
- TO\_SINGLE\_BYTE function, 4-104
- top-N queries, 2-61
- transaction control statements, 6-4
  - PL/SQL support of, 6-4
- TRANSACTION\_AUDITING parameter
  - of ALTER SYSTEM, 7-118
- transactions
  - allowing to complete, 7-106
  - assigning
    - rollback segment to, 7-602
    - to a specific rollback segment, 7-603
  - automatically committing, 7-230
  - commenting on, 7-231
  - distributed, forcing, 7-84
  - ending, 7-230
  - implicit commit of, 6-2, 6-4, 6-5
  - in-doubt
    - committing, 7-230
    - forcing, 7-231
  - isolation level, 7-602
  - locks, releasing, 7-230
  - read-committed isolation mode, 7-603
  - read-only, 7-602
  - read-write, 7-602
  - rolling back, 7-107, 7-367, 7-564
    - to a savepoint, 7-565
  - savepoints for, 7-567
  - serializable isolation mode, 7-603
- TRANSLATE ... USING function, 4-106
- TRANSLATE function, 4-105
- triggers
  - AFTER, 7-428
  - BEFORE, 7-428
  - compiling, 7-183
  - creating, 7-427
  - creating multiple, 7-429
  - database
    - altering, 7-183
    - dropping, 7-505, 7-510
  - disabling, 7-165, 7-183
  - enabling, 7-165, 7-183, 7-427
  - executing
    - with a PL/SQL block, 7-433
    - with an external procedure, 7-433
  - granting
    - system privileges on, 7-527
  - INSTEAD OF, 7-429
    - dropping, 7-458
  - on database events, 7-431
  - on DDL events, 7-430
  - on DML operations, 7-429
  - on views, 7-429
  - order of firing, 7-429
  - re-creating, 7-428
  - removing from the database, 7-505
  - restrictions on, 7-433
  - row values
    - old and new, 7-432
  - row, specifying, 7-432
  - SQL examples, 7-434
  - statement, 7-432
- TRIM function, 4-107
- TRUNC function
  - date function, 4-109
  - format models, 4-117
  - number function, 4-108
- TRUNCATE PARTITION clause
  - of ALTER TABLE, 7-160
- TRUNCATE statement, 7-611
- TRUNCATE SUBPARTITION clause
  - of ALTER TABLE, 7-160
- TRUST parameter
  - of PRAGMA RESTRICT\_REFERENCES, 7-189
- Trusted Oracle, 1-5
- type constructor expressions, 5-6
- TYPES clause
  - of ASSOCIATE STATISTICS, 7-210, 7-212
- types. See object types or datatypes

## U

---

- U number format element, 2-41
- UID function, 4-109
- unary operators, 3-1
- UNION ALL set operator, 3-13, 7-579
- UNION set operator, 3-13, 7-579
- UNIQUE clause
  - of constraint\_clause, 7-238
  - of CREATE INDEX, 7-296
  - of CREATE TABLE, 7-393
  - of SELECT, 7-573
- unique constraints
  - enabling, 7-407
  - index on, 7-407
- unique indexes, 7-296
- unique queries, 7-573
- universal rowids. *See* urowids
- UNLIMITED TABLESPACE system
  - privilege, 7-527
- unnesting collections, 7-576
  - examples, 7-593
- UNRECOVERABLE, 7-37, 7-396
  - See also* NOLOGGING clause
- unsorted indexes, 7-300
- UNUSABLE clause
  - of ALTER INDEX, 7-40
- UNUSABLE LOCAL INDEXES clause
  - of ALTER MATERIALIZED VIEW, 7-51
  - of ALTER TABLE, 7-157
- UPDATE ANY TABLE system privilege, 7-527
- UPDATE object privilege, 7-535
- UPDATE statement, 7-615
  - triggers on, 7-429
- UPPER function, 4-110
- UROWID datatype, 2-27
- urowids
  - and foreign tables, 2-27
  - and heap-organized tables, 2-28
  - and index-organized tables, 2-27
  - description of, 2-27
- USE ROLLBACK SEGMENT clause
  - of SET TRANSACTION, 7-603
- USE\_CONCAT hint, 2-65
- USE\_HASH hint, 2-66
- USE\_MERGE hint, 2-66
- USE\_NL hint, 2-66
- USE\_STORED\_OUTLINES parameter
  - of ALTER SESSION, 7-95
  - of ALTER SYSTEM, 7-118
- USER function, 4-110
- USER\_COL\_COMMENTS view, 7-228
- USER\_DUMP\_DEST parameter
  - of ALTER SYSTEM, 7-119
- USER\_TAB\_COMMENTS view, 7-228
- user-defined functions, 4-118
  - name precedence of, 4-119
  - naming conventions, 4-120
  - restrictions on, 7-286
- user-defined operators, 3-16
- user-defined statistics
  - dropping, 7-478, 7-480, 7-482, 7-492, 7-501, 7-506
- user-defined types
  - categories of, 2-30
  - defining, 7-440
- USERENV function, 4-111
- users
  - allocating space for, 7-453
  - assigning
    - default roles, 7-195
    - profiles, 7-453
  - authenticating to a remote server, 7-275
  - changing global authentication, 7-195
  - creating, 7-451
  - default tablespaces of, 7-453
  - denying access to tables and views, 7-547
  - external, 7-366, 7-452
  - global, 7-366, 7-452
  - granting
    - system privileges on, 7-528
  - local, 7-366, 7-452
  - locking accounts of, 7-453
  - maximum concurrent, 7-112
  - password expiration of, 7-453
  - removing from the database, 7-509
  - SQL examples, 7-453
  - temporary tablespaces for, 7-453
- USING BFILE clause
  - of CREATE JAVA, 7-314
- USING BLOB clause



- of CREATE JAVA, 7-314
- USING clause
  - of ASSOCIATE STATISTICS, 7-211, 7-212
  - of CREATE DATABASE LINK, 7-275
  - of CREATE INDEXTYPE, 7-310
  - of CREATE OPERATOR, 7-339, 7-341
- USING CLOB clause
  - of CREATE JAVA, 7-314
- USING INDEX clause
  - of ALTER MATERIALIZED VIEW, 7-52
  - of constraint\_clause, 7-244
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT, 7-324
  - of CREATE TABLE, 7-389, 7-407
- USING ROLLBACK SEGMENT clause
  - of ALTER MATERIALIZED  
VIEW...REFRESH, 7-53
  - of CREATE MATERIALIZED VIEW/  
SNAPSHOT...REFRESH, 7-326
- UTLCHN.SQL script, 7-207
- UTLEXPT1.SQL script, 7-163, 7-245, 7-407
- UTLXPLAN.SQL script, 7-512

## V

---

- V number format element, 2-41
- VSNLS\_PARAMETERS view
- VALIDATE REF UPDATE clause
  - of ANALYZE, 7-206
- VALIDATE STRUCTURE clause
  - of ANALYZE, 7-206
- VALUE function, 4-112
- VALUES clause
  - of CREATE INDEX, 7-302
  - of INSERT, 7-542
- VALUES LESS THAN clause
  - of CREATE TABLE, 7-403
- VARCHAR datatype, 2-17
  - DB2, 2-29
  - SQL/DS, 2-29
- VARCHAR2 datatype, 2-16
  - converting to NUMBER, 2-40
- VARGRAPHIC datatype (SQL/DS or DB2), 2-30
- variable expressions, 5-4
- VARIANCE function, 4-115

- VARP function, 4-113
- VARRAY storage clause
  - of ALTER TABLE, 7-142
  - of CREATE TABLE, 7-385, 7-399
- varrays, 2-31
  - changing returned value, 7-145
  - compared with nested tables, 2-35
  - comparison rules, 2-35
  - creating, 7-439, 7-444
  - dropping the body of, 7-508
  - dropping the specification of, 7-506
  - storage characteristics of, 7-142, 7-146, 7-399
  - storing out of line, 2-31
- VARS function, 4-114
- varying arrays. *See* varrays
- views
  - adding rows to the base table of, 7-540
  - changing
    - definition, 7-511
    - values in base tables, 7-616
  - creating
    - before base tables, 7-458
    - comments about, 7-228
    - multiple, 7-369
  - defining, 7-456
  - granting
    - system privileges on, 7-529
  - recompiling, 7-198
  - re-creating, 7-458
  - remote, accessing, 7-273
  - removing
    - from the database, 7-511
    - rows from the base table of, 7-465
  - renaming, 7-554
  - retrieving data from, 7-572
  - subquery of, 7-459
    - restricting, 7-461
  - synonyms for, 7-377
- VSIZE function, 4-116

## W

---

- W date format element, 2-46
- WHEN clause
  - of CREATE TRIGGER, 7-433

WHENEVER NOT SUCCESSFUL clause  
of NOAUDIT schema\_objects, 7-553

WHENEVER SUCCESSFUL clause  
of AUDIT schema\_objects, 7-222  
of AUDIT sql\_statements, 7-215  
of NOAUDIT schema\_objects, 7-553  
of NOAUDIT sql\_statements, 7-551

WHERE clause  
of DELETE, 7-467  
of SELECT, 5-20, 7-577  
of UPDATE, 7-620

WITH ADMIN OPTION clause  
of GRANT system\_privileges\_and\_roles, 7-521

WITH CHECK OPTION clause  
of CREATE VIEW, 7-456, 7-461  
of DELETE, 7-467  
of INSERT, 7-542  
of SELECT, 7-570, 7-575  
of UPDATE, 7-618

WITH GRANT OPTION clause  
of GRANT object\_privileges, 7-534

WITH INDEX CONTEXT clause  
of CREATE OPERATOR, 7-339, 7-341

WITH OBJECT IDENTIFIER clause  
of CREATE VIEW, 7-459

WITH OBJECT OID. *See* WITH OBJECT IDENTIFIER.

WITH PRIMARY KEY clause  
of ALTER MATERIALIZED VIEW, 7-53  
of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336  
of CREATE MATERIALIZED VIEW/  
SNAPSHOT...REFRESH, 7-324

WITH READ ONLY clause  
of CREATE VIEW, 7-456, 7-461  
of DELETE, 7-467  
of INSERT, 7-542  
of SELECT, 7-570, 7-575  
of UPDATE, 7-618

WITH ROWID clause  
of column ref constraints, 7-242  
of CREATE MATERIALIZED VIEW LOG/  
SNAPSHOT LOG, 7-336  
of CREATE MATERIALIZED VIEW/  
SNAPSHOT...REFRESH, 7-324

WNDS parameter  
of PRAGMA RESTRICT\_REFERENCES, 7-189

WNPS parameter  
of PRAGMA RESTRICT\_REFERENCES, 7-189

WW date format element, 2-46

---

**X**

---

X number format element, 2-41

---

**Y**

---

Y date format element, 2-46

YEAR date format element, 2-46

YY date format element, 2-46

YYY date format element, 2-46

YYYY date format element, 2-46