

## **Oracle® Database**

Application Developer's Guide - Object-Relational Features

10g Release 1 (10.1)

**Part No. B10799-01**

December 2003

Oracle Database Application Developer's Guide - Object-Relational Features, 10g Release 1 (10.1)

Part No. B10799-01

Copyright © 1996, 2003 Oracle Corporation. All rights reserved.

Contributors: Geeta Arora, Eric Belden, Chandrasekharan Iyer, Geoff Lee, Anand Manikutty, Valarie Moore, Magdi Morsi, Helen Yeh, Adiel Yoaz, Qin Yu

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, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. 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, and Oracle9i, Oracle Store, PL/SQL, Pro\*C, Pro\*C/C++, and SQL\*Plus are trademarks or registered trademarks of Oracle Corporation. Other names may be trademarks of their respective owners.

---

---

# Contents

<b>Send Us Your Comments .....</b>	<b>xiii</b>
<b>Preface.....</b>	<b>xv</b>
Audience .....	xvi
Organization.....	xvi
Related Documentation .....	xvii
Conventions.....	xviii
Documentation Accessibility .....	xxi
<b>What's New in Object-Relational Features? .....</b>	<b>xxiii</b>
Oracle 10g Release 1 (10.1) New in Object-Relational Features.....	xxiv
Oracle9i Release 2 (9.2) New in Object-Relational Features .....	xxiv
Oracle9i Release 1 (9.0.1) New in Object-Relational Features.....	xxv
<b>1 Introduction to Oracle Objects</b>	
<b>About Oracle Objects.....</b>	<b>1-2</b>
<b>Advantages of Objects.....</b>	<b>1-2</b>
<b>Key Features of the Object-Relational Model.....</b>	<b>1-3</b>
Core Database Key Features .....	1-4
Object Types.....	1-4
Objects.....	1-5
Object Methods.....	1-6
Type Inheritance.....	1-7
Type Evolution .....	1-7

Object Tables .....	1-8
Row Objects and Column Objects.....	1-9
Object Views.....	1-9
References .....	1-9
Collections .....	1-12
Language Binding Features.....	1-13

## 2 Basic Components of Oracle Objects

<b>SQL Object Types and References</b> .....	2-2
Null Objects and Attributes .....	2-2
Character Length Semantics .....	2-3
Constraints for Object Tables.....	2-4
Indexes for Object Tables.....	2-5
Triggers for Object Tables.....	2-5
Rules for REF Columns and Attributes.....	2-6
Name Resolution .....	2-6
When Table Aliases Are Required.....	2-7
Restriction on Using User-Defined Types with a Remote Database.....	2-8
<b>Object Methods</b> .....	2-8
Member Methods.....	2-9
Methods for Comparing Objects .....	2-10
Map Methods .....	2-10
Order Methods.....	2-12
Guidelines for Comparison Methods .....	2-13
Comparison Methods in Type Hierarchies.....	2-13
Static Methods.....	2-13
Constructor Methods .....	2-14
External Implemented Methods.....	2-15
<b>Inheritance in SQL Object Types</b> .....	2-15
Types and Subtypes.....	2-15
FINAL and NOT FINAL Types and Methods.....	2-17
Creating Subtypes.....	2-18
NOT INSTANTIABLE Types and Methods .....	2-19
Inheriting, Overloading, and Overriding Methods.....	2-20
Overloading Methods .....	2-21

Redefining Methods.....	2-21
Restrictions on Overriding Methods .....	2-22
Dynamic Method Dispatch .....	2-23
Substituting Types in a Type Hierarchy .....	2-24
Column and Row Substitutability.....	2-25
Using OBJECT_VALUE and OBJECT_ID with Substitutable Rows.....	2-26
Subtypes Having Supertype Attributes .....	2-27
REF Columns and Attributes.....	2-28
Collection Elements.....	2-28
Creating Subtypes After Creating Substitutable Columns.....	2-28
Dropping Subtypes After Creating Substitutable Columns.....	2-29
Turning Off Substitutability in a New Table .....	2-29
Constraining Substitutability.....	2-30
Modifying Substitutability .....	2-31
Restrictions on Modifying Substitutability.....	2-31
Assignments Across Types .....	2-32
Objects and REFs to Objects.....	2-32
Collection Assignments.....	2-34
Comparisons of Objects, REF Variables, and Collections.....	2-34
Comparing Object Instances .....	2-34
Comparing REF Variables.....	2-35
<b>Functions and Operators Useful with Objects.....</b>	<b>2-35</b>
CAST.....	2-35
CURSOR.....	2-36
DEREF .....	2-36
IS OF <i>type</i> .....	2-36
REF.....	2-38
SYS_TYPEID.....	2-38
TABLE().....	2-39
TREAT .....	2-39
VALUE.....	2-41

### 3 Support for Collection Datatypes

<b>Creating Collection Datatypes .....</b>	<b>3-2</b>
Creating an Instance of a VARRAY or Nested Table.....	3-2

Constructor Methods for Collections.....	3-2
Varrays .....	3-3
Nested Tables .....	3-4
Specifying a Tablespace When Storing a Nested Table .....	3-6
Varray Storage.....	3-6
Increasing the Size and Precision of VARRAYs and Nested Tables .....	3-7
Increasing VARRAY Limit Size .....	3-8
Creating a Varray Containing LOB References.....	3-8
Multilevel Collection Types .....	3-9
Nested Table Storage Tables for Multilevel Collection Types.....	3-9
Assignment and Comparison of Multilevel Collections.....	3-11
Constructors for Multilevel Collections .....	3-11
<b>Operations on Collection Datatypes .....</b>	<b>3-12</b>
Querying Collections.....	3-12
Nesting Results of Collection Queries.....	3-12
Unnesting Results of Collection Queries .....	3-13
Unnesting Queries Containing Table Expression Subqueries .....	3-14
Unnesting Queries with Multilevel Collections.....	3-15
Performing DML Operations on Collections.....	3-15
Performing DML on Multilevel Collections .....	3-16
Comparisons of Collections .....	3-17
Equal and Not Equal Comparisons .....	3-17
In Comparisons.....	3-18
Subset of Multiset Comparison .....	3-18
Member of a Nested Table Comparison .....	3-18
Empty Comparison .....	3-19
Set Comparison.....	3-19
Multisets Operations .....	3-19
CARDINALITY.....	3-19
COLLECT.....	3-20
MULTISET EXCEPT.....	3-20
MULTISET INTERSECTION .....	3-20
MULTISET UNION.....	3-21
POWERMULTISET .....	3-22
POWERMULTISET_BY_CARDINALITY.....	3-22

SET .....	3-22
-----------	------

## 4 Object Support in Oracle Programming Environments

<b>SQL</b> .....	4-2
<b>PL/SQL</b> .....	4-2
<b>Oracle Call Interface (OCI)</b> .....	4-2
Associative Access in OCI Programs .....	4-3
Navigational Access in OCI Programs .....	4-4
Object Cache .....	4-4
Building an OCI Program That Manipulates Objects .....	4-5
Defining User-Defined Constructors in C .....	4-6
<b>Pro*C/C++</b> .....	4-6
Associative Access in Pro*C/C++ .....	4-7
Navigational Access in Pro*C/C++ .....	4-7
Converting Between Oracle Types and C Types .....	4-8
Oracle Type Translator (OTT) .....	4-8
<b>Oracle C++ Call Interface (OCCI)</b> .....	4-8
OCCI Associative Relational and Object Interfaces .....	4-9
The OCCI Navigational Interface .....	4-9
<b>Oracle Objects For OLE (OO4O)</b> .....	4-10
Representing Objects in Visual Basic (OraObject) .....	4-11
Representing REFs in Visual Basic (OraRef) .....	4-12
Representing VARRAYs and Nested Tables in Visual Basic (OraCollection) .....	4-12
<b>Java: JDBC, Oracle SQLJ, JPublisher, and SQLJ Object Types</b> .....	4-12
JDBC Access to Oracle Object Data .....	4-13
SQLJ Access to Oracle Object Data .....	4-14
Choosing a Data Mapping Strategy .....	4-14
Using JPublisher to Create Java Classes for JDBC and SQLJ Programs .....	4-14
What JPublisher Produces for a User-Defined Object Type .....	4-15
Java Object Storage .....	4-16
Representing SQLJ Types to the Server .....	4-17
Creating SQLJ Object Types .....	4-17
Additional Notes About Mapping .....	4-18
Evolving SQLJ Types .....	4-19
Constraints .....	4-19

Querying SQLJ Objects .....	4-20
Inserting Java Objects.....	4-20
Updating SQLJ Objects .....	4-20
Defining User-Defined Constructors in Java .....	4-20
<b>XML</b> .....	4-21

## **5 Applying an Object Model to Relational Data**

<b>Why Use Object Views</b> .....	5-2
<b>Defining Object Views</b> .....	5-3
<b>Using Object Views in Applications</b> .....	5-4
<b>Nesting Objects in Object Views</b> .....	5-4
<b>Identifying Null Objects in Object Views</b> .....	5-6
<b>Using Nested Tables and Varrays in Object Views</b> .....	5-6
Single-Level Collections in Object Views.....	5-6
Multilevel Collections in Object Views.....	5-7
<b>Specifying Object Identifiers for Object Views</b> .....	5-9
<b>Creating References to View Objects</b> .....	5-10
<b>Modelling Inverse Relationships with Object Views</b> .....	5-11
<b>Updating Object Views</b> .....	5-12
Updating Nested Table Columns in Views .....	5-13
Using INSTEAD OF Triggers to Control Mutating and Validation.....	5-13
<b>Applying the Object Model to Remote Tables</b> .....	5-14
<b>Defining Complex Relationships in Object Views</b> .....	5-15
Tables and Types to Demonstrate Circular View References .....	5-16
Creating Object Views with Circular References .....	5-17
<b>Object View Hierarchies</b> .....	5-19
Creating an Object View Hierarchy .....	5-21
The Flat Model .....	5-21
The Horizontal Model.....	5-23
The Vertical Model .....	5-25
Querying a View in a Hierarchy.....	5-26
Privileges for Operations on View Hierarchies.....	5-27

## **6 Managing Oracle Objects**

<b>Privileges on Object Types and Their Methods</b> .....	6-2
---	-----

System Privileges for Object Types.....	6-2
Schema Object Privileges.....	6-2
Using Types in New Types or Tables.....	6-3
Example: Privileges on Object Types.....	6-3
Privileges on Type Access and Object Access.....	6-4
<b>Dependencies and Incomplete Types</b> .....	6-5
Completing Incomplete Types .....	6-7
Manually Recompiling a Type.....	6-8
Type Dependencies of Substitutable Tables and Columns .....	6-8
The FORCE Option .....	6-9
<b>Synonyms for User-Defined Types</b> .....	6-9
Creating a Type Synonym.....	6-9
Using a Type Synonym.....	6-10
Describing Schema Objects That Use Synonyms.....	6-10
Dependents of Type Synonyms .....	6-11
Restriction on Replacing a Type Synonym.....	6-11
Dropping Type Synonyms.....	6-11
Renaming Type Synonyms .....	6-12
Public Type Synonyms and Local Schema Objects .....	6-12
<b>Performance Tuning</b> .....	6-12
<b>Tools Providing Support for Objects</b> .....	6-13
<b>Utilities Providing Support for Objects</b> .....	6-14

## 7 Advanced Topics for Oracle Objects

<b>Storage of Objects</b> .....	7-2
Leaf-Level Attributes .....	7-2
How Row Objects Are Split Across Columns .....	7-2
Hidden Columns for Tables with Column Objects .....	7-2
Hidden Columns for Substitutable Columns and Tables.....	7-3
REFs .....	7-4
Internal Layout of Nested Tables .....	7-4
Internal Layout of VARRAYs .....	7-5
<b>Creating Indexes on Typeids or Attributes</b> .....	7-5
Indexing a Type Discriminant Column.....	7-5
Indexing Subtype Attributes of a Substitutable Column.....	7-6

<b>Type Evolution</b> .....	7-7
Changes Involved When a Type Is Altered .....	7-9
Steps to Change a Type.....	7-10
Validating a Type.....	7-11
If a Type Change Validation Fails .....	7-15
ALTER TYPE Statement for Type Evolution .....	7-15
ALTER TABLE Statement for Type Evolution .....	7-19
The Attribute-Value Constructor .....	7-19
Constructors and Type Evolution .....	7-19
Advantages of User-Defined Constructors.....	7-20
Defining and Implementing User-Defined Constructors .....	7-20
Overloading and Hiding Constructors.....	7-22
Calling User-Defined Constructors.....	7-22
Constructors for SQLJ Object Types .....	7-24
<b>Transient and Generic Types</b> .....	7-25
<b>User-Defined Aggregate Functions</b> .....	7-28
<b>Partitioning Tables That Contain Oracle Objects</b> .....	7-29
How Locators Improve the Performance of Nested Tables .....	7-30

## 8 Design Considerations for Oracle Objects

<b>General Storage Considerations for Objects</b> .....	8-2
Storing Objects as Columns or Rows.....	8-2
Column Object Storage .....	8-2
Row Object Storage in Object Tables .....	8-5
Storage Considerations for Object Identifiers (OIDs).....	8-6
Primary-Key Based OIDs .....	8-6
<b>Performance of Object Comparisons</b> .....	8-7
<b>Design Considerations for REFs</b> .....	8-8
Storage Size of REFs .....	8-8
Integrity Constraints for REF Columns.....	8-8
Performance and Storage Considerations for Scoped REFs.....	8-9
Indexing Scoped REFs .....	8-9
Speeding up Object Access Using the WITH ROWID Option.....	8-10
<b>Design Considerations for Collections</b> .....	8-11
Viewing Object Data in Relational Form with Unnesting Queries .....	8-11

Using Procedures and Functions in Unnesting Queries.....	8-12
Storage Considerations for Varrays.....	8-13
Propagating VARRAY Size Change .....	8-13
Performance of Varrays Versus Nested Tables.....	8-14
Design Considerations for Nested Tables.....	8-14
Nested Table Storage .....	8-14
Nested Table Indexes.....	8-17
Nested Table Locators .....	8-18
Optimizing Set Membership Queries.....	8-19
Design Considerations for Multilevel Collections.....	8-19
<b>Design Considerations for Methods.....</b>	<b>8-24</b>
Choosing a Language for Method Functions .....	8-25
Static Methods.....	8-27
Using SELF IN OUT NOCOPY with Member Procedures .....	8-28
Function-Based Indexes on the Return Values of Type Methods .....	8-28
<b>Writing Reusable Code Using Invoker Rights.....</b>	<b>8-29</b>
<b>Replicating Object Tables and Columns.....</b>	<b>8-31</b>
Replicating Columns of Object, Collection, or REF Type.....	8-31
Replicating Object Tables .....	8-32
<b>Constraints on Objects.....</b>	<b>8-33</b>
<b>Considerations Related to Type Evolution.....</b>	<b>8-33</b>
Pushing a Type Change Out to Clients .....	8-33
Changing Default Constructors.....	8-34
Altering the FINAL Property of a Type .....	8-34
<b>Parallel Queries with Oracle Objects .....</b>	<b>8-34</b>
<b>Design Consideration Tips and Techniques .....</b>	<b>8-35</b>
Deciding Whether to Evolve a Type or Create a Subtype Instead .....	8-35
How ANYDATA Differs from User-Defined Types .....	8-36
Polymorphic Views: An Alternative to an Object View Hierarchy .....	8-36
The SQLJ Object Type .....	8-37
The Intended Use of SQLJ Object Types.....	8-37
Actions Performed When Creating a SQLJ Object Type .....	8-38
Uses of SQLJ Object Types.....	8-38
Uses of Custom Object Types.....	8-38
Differences Between SQLJ and Custom Object Types Through JDBC.....	8-39

Miscellaneous Tips .....	8-39
Column Substitutability and the Number of Attributes in a Hierarchy .....	8-39
Circular Dependencies Among Types.....	8-40

## **9 A Sample Application Using Object-Relational Features**

<b>Introduction to the Sample Application.....</b>	<b>9-2</b>
--	------------

## **Index**

---

---

# Send Us Your Comments

## Oracle Database Application Developer's Guide - Object-Relational Features, 10g Release 1 (10.1)

Part No. B10799-01

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 title and part number of the documentation and the chapter, section, and page number (if available). You can send comments to us in the following ways:

- Electronic mail: [infodev\\_us@oracle.com](mailto:infodev_us@oracle.com)
- FAX: (650) 506-7227 Attn: Server Technologies Documentation Manager
- Postal service:  
Oracle Corporation  
Server Technologies Documentation  
500 Oracle Parkway, Mailstop 4op11  
Redwood Shores, CA 94065  
USA

If you would like a reply, please give your name, address, telephone number, and electronic mail address (optional).

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



---

# Preface

*Oracle Database Application Developer's Guide - Object-Relational Features* describes how to use the object-relational features of the Oracle Server, 10g Release 1 (10.1). Information in this guide applies to versions of the Oracle Server that run on all platforms, and does not include system-specific information.

This preface contains these topics:

- [Audience](#)
- [Organization](#)
- [Related Documentation](#)
- [Conventions](#)
- [Documentation Accessibility](#)

## Audience

*Oracle Database Application Developer's Guide - Object-Relational Features* is intended for programmers developing new applications or converting existing applications to run in the Oracle environment. The object-relational features are often used in content management, data warehousing, data/information integration, and similar applications that deal with complex structured data. The object views feature can be valuable when writing new C++, Java, or XML applications on top of an existing relational schema.

This guide assumes that you have a working knowledge of application programming and that you are familiar with the use of Structured Query Language (SQL) to access information in relational database systems.

## Organization

This document contains:

### **Chapter 1, "Introduction to Oracle Objects"**

Introduces the key features and explains the advantages of the object-relational model.

### **Chapter 2, "Basic Components of Oracle Objects"**

Explains the basic concepts and terminology that you need to work with Oracle Objects.

### **Chapter 3, "Support for Collection Datatypes"**

Discusses collection datatypes and operations on collection datatypes.

### **Chapter 4, "Object Support in Oracle Programming Environments"**

Summarizes the object-relational features in SQL and PL/SQL; Oracle Call Interface (OCI); Pro\*C/C++; Oracle Objects For OLE; and Java, JDBC, and Oracle SQLJ. The information in this chapter is high-level, for education and planning. The following chapters explain how to use the object-relational features in greater detail.

### **Chapter 5, "Applying an Object Model to Relational Data"**

Explains object views, which allow you to develop object-oriented applications without changing the underlying relational schema.

### **Chapter 6, "Managing Oracle Objects"**

Explains how to perform essential operations with objects and object types.

### **Chapter 7, "Advanced Topics for Oracle Objects"**

Discusses features that you might need to manage storage and performance as you scale up an object-oriented application.

### **Chapter 8, "Design Considerations for Oracle Objects"**

Explains the implementation and performance characteristics of the Oracle object-relational model.

### **Chapter 9, "A Sample Application Using Object-Relational Features"**

Demonstrates how a relational program can be rewritten as an object-oriented one, schema and all.

## **Related Documentation**

For more information, see these Oracle resources:

- *PL/SQL User's Guide and Reference* for information on PL/SQL, the procedural language extension to Oracle SQL
- *Oracle Database Application Developer's Guide - Fundamentals* for general information about developing applications
- *Oracle XML DB Developer's Guide* and *Oracle XML Developer's Kit Programmer's Guide* for information about developing applications with XML
- *Oracle Database JDBC Developer's Guide and Reference* and *Oracle Database Java Developer's Guide* to use Oracle object-relational features through Java
- *Oracle Call Interface Programmer's Guide* and *Oracle C++ Call Interface Programmer's Guide* for information on using the Oracle Call Interface (OCI) and Oracle C++ Call Interface to build third-generation language (3GL) applications that access the Oracle Server
- *Pro\*C/C++ Programmer's Guide* for information on Oracle's Pro\* series of precompilers, which allow you to embed SQL and PL/SQL in 3GL application programs written in Ada, C, C++, COBOL, or FORTRAN
- *Oracle Database SQL Reference* and *Oracle Database Administrator's Guide* for information on SQL
- *Oracle Database Concepts* for information on basic Oracle concepts

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

Printed documentation is available for sale in the Oracle Store at

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

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

<http://otn.oracle.com/membership/>

If you already have a username and password for OTN, then you can go directly to the documentation section of the OTN Web site at

<http://otn.oracle.com/documentation/>

## Conventions

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

- [Conventions in Text](#)
- [Conventions in Code Examples](#)

### Conventions in Text

We use various conventions in text to help you more quickly identify special terms. The following table describes those conventions and provides examples of their use.

Convention	Meaning	Example
<b>Bold</b>	Bold typeface indicates terms that are defined in the text or terms that appear in a glossary, or both.	When you specify this clause, you create an <b>index-organized table</b> .
<i>Italics</i>	Italic typeface indicates book titles or emphasis.	<i>Oracle Database Concepts</i> Ensure that the recovery catalog and target database do <i>not</i> reside on the same disk.

Convention	Meaning	Example
UPPERCASE monospace (fixed-width) font	Uppercase monospace typeface indicates elements supplied by the system. Such elements include parameters, privileges, datatypes, RMAN keywords, SQL keywords, SQL*Plus or utility commands, packages and methods, as well as system-supplied column names, database objects and structures, usernames, and roles.	You can specify this clause only for a NUMBER column.  You can back up the database by using the BACKUP command.  Query the TABLE_NAME column in the USER_TABLES data dictionary view.  Use the DBMS_STATS.GENERATE_STATS procedure.
lowercase monospace (fixed-width) font	Lowercase monospace typeface indicates executables, filenames, directory names, and sample user-supplied elements. Such elements include computer and database names, net service names, and connect identifiers, as well as user-supplied database objects and structures, column names, packages and classes, usernames and roles, program units, and parameter values.  <b>Note:</b> Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	Enter sqlplus to open SQL*Plus.  The password is specified in the orapwd file.  Back up the datafiles and control files in the /disk1/oracle/dbs directory.  The department_id, department_name, and location_id columns are in the hr.departments table.  Set the QUERY_REWRITE_ENABLED initialization parameter to true.  Connect as oe user.  The JRepuTil class implements these methods.
lowercase italic monospace (fixed-width) font	Lowercase italic monospace font represents placeholders or variables.	You can specify the <i>parallel_clause</i> .  Run <i>Uold_release</i> .SQL where <i>old_release</i> refers to the release you installed prior to upgrading.

## Conventions in Code Examples

Code examples illustrate SQL, PL/SQL, SQL\*Plus, or other command-line statements. They are displayed in a monospace (fixed-width) font and separated from normal text as shown in this example:

```
SELECT username FROM DBA_USERS WHERE username = 'MIGRATE';
```

The following table describes typographic conventions used in code examples and provides examples of their use.

Convention	Meaning	Example
[ ]	Brackets enclose one or more optional items. Do not enter the brackets.	DECIMAL ( <i>digits</i> [ , <i>precision</i> ])
{ }	Braces enclose two or more items, one of which is required. Do not enter the braces.	{ENABLE   DISABLE}
	A vertical bar represents a choice of two or more options within brackets or braces. Enter one of the options. Do not enter the vertical bar.	{ENABLE   DISABLE} [COMPRESS   NOCOMPRESS]
...	Horizontal ellipsis points indicate either: <ul style="list-style-type: none"> <li>That we have omitted parts of the code that are not directly related to the example</li> <li>That you can repeat a portion of the code</li> </ul>	CREATE TABLE ... AS <i>subquery</i> ;  SELECT <i>col1</i> , <i>col2</i> , ... , <i>coln</i> FROM employees;
.	Vertical ellipsis points indicate that we have omitted several lines of code not directly related to the example.	SQL> SELECT NAME FROM V\$DATAFILE; NAME ----- /fsl/dbs/tbs_01.dbf /fsl/dbs/tbs_02.dbf . . . /fsl/dbs/tbs_09.dbf 9 rows selected.
Other notation	You must enter symbols other than brackets, braces, vertical bars, and ellipsis points as shown.	acctbal NUMBER(11,2); acct CONSTANT NUMBER(4) := 3;
<i>Italics</i>	Italicized text indicates placeholders or variables for which you must supply particular values.	CONNECT SYSTEM/ <i>system_password</i>  DB_NAME = <i>database_name</i>
UPPERCASE	Uppercase typeface indicates elements supplied by the system. We show these terms in uppercase in order to distinguish them from terms you define. Unless terms appear in brackets, enter them in the order and with the spelling shown. However, because these terms are not case sensitive, you can enter them in lowercase.	SELECT last_name, employee_id FROM employees;  SELECT * FROM USER_TABLES;  DROP TABLE hr.employees;

Convention	Meaning	Example
lowercase	<p>Lowercase typeface indicates programmatic elements that you supply. For example, lowercase indicates names of tables, columns, or files.</p> <p><b>Note:</b> Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.</p>	<pre>SELECT last_name, employee_id FROM employees;  sqlplus hr/hr  CREATE USER mjones IDENTIFIED BY ty3MU9;</pre>

## Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For additional information, visit the Oracle Accessibility Program Web site at

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

**Accessibility of Code Examples in Documentation** JAWS, a Windows screen reader, may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, JAWS may not always read a line of text that consists solely of a bracket or brace.

**Accessibility of Links to External Web Sites in Documentation** This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.



---

---

# What's New in Object-Relational Features?

This section describes the new object-relational features of Oracle 10g Release 1 (10.1). New features information from previous releases is also retained to help those users upgrading to the current release.

The following sections describe the new features in Oracle object-relational features:

- [Oracle 10g Release 1 \(10.1\) New in Object-Relational Features](#)
- [Oracle9i Release 2 \(9.2\) New in Object-Relational Features](#)
- [Oracle9i Release 1 \(9.0.1\) New in Object-Relational Features](#)

## Oracle 10g Release 1 (10.1) New in Object-Relational Features

New object-relational features for Oracle 10g Release 1 (10.1) include the following collection enhancements:

- New functionality for nested table and varray storage, including the evolution of varray size and specification of a tablespace when storing nested tables. See "[Creating Collection Datatypes](#)" on page 3-2.
- New functionality for nested table comparisons and ANSI SQL multiset operations for nested tables. See "[Operations on Collection Datatypes](#)" on page 3-12.

## Oracle9i Release 2 (9.2) New in Object-Relational Features

New object-relational features for Oracle9i release 2 (9.2) include:

- Character length semantics  
Lengths for character types CHAR and VARCHAR2 may be specified as a number of characters, instead of bytes, in object attributes and collections even if some of the characters consist of multiple bytes. See "[Character Length Semantics](#)" on page 2-3.
- Modifying substitutability  
In an existing table, you can change the substitutability of an object column by using an ALTER TABLE statement with the [NOT] SUBSTITUTABLE AT ALL LEVELS clause. See "[Modifying Substitutability](#)" on page 2-31.
- Type synonyms  
Synonyms can be defined for user-defined types so that a type can be used without having to qualify its name with the name of the schema in which the type was defined. See "[Synonyms for User-Defined Types](#)" on page 6-9.
- User-defined constructors  
User-defined constructor functions make possible custom initialization of newly created object instances. They also make it possible to evolve a type without having to update calls to constructors in existing code to accommodate a newly added attribute. See "[Advantages of User-Defined Constructors](#)" on page 7-20.

# Oracle9i Release 1 (9.0.1) New in Object-Relational Features

New object-relational features for Oracle9i release 1 (9.0.1) include:

- SQL type inheritance  
User-defined types can be specialized as subtypes in a SQL type hierarchy. See ["Inheritance in SQL Object Types"](#) on page 2-15.
- Object view hierarchies  
Hierarchies can be created of object views based on some or all of the types in a type hierarchy. Object view hierarchies simplify targeting a particular subtype (and perhaps its subtypes) in queries and other operations. See ["Object View Hierarchies"](#) on page 5-19.
- Type evolution  
User-defined SQL types and their independent objects can be changed, or evolved, instead of having to be re-created. See ["Type Evolution"](#) on page 7-7.
- User-defined aggregate functions  
Custom aggregate functions can be defined for working with complex data. See ["User-Defined Aggregate Functions"](#) on page 7-28.
- Generic and transient datatypes  
External procedures can be given fields or parameters of a generic type, such as `SYS.ANYTYPE`, that can contain values of any scalar or user-defined type, making it unnecessary to implement multiple versions of the same external procedure just to handle multiple datatypes. See ["Transient and Generic Types"](#) on page 7-25.
- Function-based indexes  
Function-based indexes can be built on type method functions. See ["Function-Based Indexes on the Return Values of Type Methods"](#) on page 8-28.
- Multilevel collections  
Collections (varrays and nested tables) can contain elements that are themselves collections or have attributes that are. See [Chapter 3, "Support for Collection Datatypes"](#) and [Chapter 5, "Applying an Object Model to Relational Data"](#).
- C++ interface to Oracle  
A C++ interface (OCCI) to Oracle, built on top of OCI, enables you to use the object-oriented features, native classes, and methods of the C++ programming

language to access the Oracle database. See "[Oracle C++ Call Interface \(OCI\)](#)" on page 4-8.

- Java object storage

You can now create SQL types mapped to existing Java classes to provide persistent storage for Java objects. Such SQL types are called SQL types of Language Java, or SQLJ types. See "[Java: JDBC, Oracle SQLJ, JPublisher, and SQLJ Object Types](#)" on page 4-12.

---

# Introduction to Oracle Objects

This chapter describes the advantages and key features of the Oracle object-relational model. The chapter contains these topics:

- [About Oracle Objects](#)
- [Advantages of Objects](#)
- [Key Features of the Object-Relational Model](#)

## About Oracle Objects

Oracle object types are user-defined types that make it possible to model real-world entities such as customers and purchase orders as objects in the database.

Oracle object technology is a layer of abstraction built on Oracle relational technology. New object types can be created from any built-in database types and any previously created object types, object references, and collection types. Metadata for user-defined types is stored in a schema that is available to SQL, PL/SQL, Java, and other published interfaces.

Object types and related object-oriented features such as variable-length arrays and nested tables provide higher-level ways to organize and access data in the database. Underneath the object layer, data is still stored in columns and tables, but you are able to work with the data in terms of the real-world entities, such as customers and purchase orders, that make the data meaningful. Instead of thinking in terms of columns and tables when you query the database, you can simply select a customer.

Internally, statements about objects are still basically statements about relational tables and columns, and you can continue to work with relational data types and store data in relational tables as before. But now you have the option to take advantage of object-oriented features too. You can begin to use object-oriented features while continuing to work with most of your data relationally, or you can go over to an object-oriented approach entirely. For instance, you can define some object data types and store the objects in columns in relational tables, which enables you to extend the system built-in types with user-defined ones. You can also create object views of existing relational data to represent and access this data according to an object model. Or you can store object data in object tables, where each row is an object.

## Advantages of Objects

In general, the object-type model is similar to the class mechanism found in C++ and Java. Like classes, objects make it easier to model complex, real-world business entities and logic, and the reusability of objects makes it possible to develop database applications faster and more efficiently. By natively supporting object types in the database, Oracle enables application developers to directly access the data structures used by their applications. No mapping layer is required between client-side objects and the relational database columns and tables that contain the data. Object abstraction and the encapsulation of object behaviors also make applications easier to understand and maintain.

Below are listed several other specific advantages that objects offer over a purely relational approach.

### **Objects Can Encapsulate Operations Along with Data**

Database tables contain only data. Objects can include the ability to perform operations that are likely to be needed on that data. Thus a purchase order object might include a method to sum the cost of all the items purchased. Or a customer object might have methods to return the customer's buying history and payment pattern. An application can simply call the methods to retrieve the information.

### **Objects Are Efficient**

Using object types makes for greater efficiency:

- Object types and their methods are stored with the data in the database, so they are available for any application to use. Developers can benefit from work that is already done and do not need to re-create similar structures in every application.
- You can fetch and manipulate a set of related objects as a single unit. A single request to fetch an object from the server can retrieve other objects that are connected to it. For example, when you select a customer object and get the customer's name, phone, and the multiple parts of his address in a single round-trip between the client and the server. When you reference a column of a SQL object type, you retrieve the whole object.

### **Objects Can Represent Part-Whole Relationships**

In a relational system, it is awkward to represent complex part-whole relationships. A piston and an engine have the same status in a table for stock items. To represent pistons as parts of engines, you must create complicated schemas of multiple tables with primary key-foreign key relationships. Object types, on the other hand, give you a rich vocabulary for describing part-whole relationships. An object can have other objects as attributes, and the attribute objects can have their own object attributes too. An entire parts-list hierarchy can be built up in this way from interlocking object types.

## **Key Features of the Object-Relational Model**

Oracle implements the object-type system as an extension of the relational model. The object-type interface continues to support standard relational database functionality such as queries (`SELECT...FROM...WHERE`), fast commits, backup and

recovery, scalable connectivity, row-level locking, read consistency, partitioned tables, parallel queries, cluster database, export and import, and loader. Plus SQL and various programmatic interfaces to Oracle; including PL/SQL, Java, Oracle Call Interface, Pro\*C/C++, and OO4O; have been enhanced with new extensions to support objects. The result is an object-relational model, which offers the intuitiveness and economy of an object interface while preserving the high concurrency and throughput of a relational database.

## Core Database Key Features

This section lists the key features and concepts of the object-relational model that are related to the database.

### Object Types

An object type is a kind of datatype. You can use it in the same ways that you use more familiar datatypes such as `NUMBER` or `VARCHAR2`. For example, you can specify an object type as the datatype of a column in a relational table, and you can declare variables of an object type. You use a variable of an object type to contain a value of that object type. A value of an object type is an instance of that type. An object instance is also called an object. [Example 1-1](#) shows how to create a user-defined object type named `person_typ`.

#### **Example 1-1** *Creating the Person\_typ Object*

```
CREATE TYPE person_typ AS OBJECT (  
    idno          NUMBER,  
    name          VARCHAR2(30),  
    phone         VARCHAR2(20),  
    MAP MEMBER FUNCTION get_idno RETURN NUMBER );  
/  
  
CREATE TYPE BODY person_typ AS  
    MAP MEMBER FUNCTION get_idno RETURN NUMBER IS  
    BEGIN  
        RETURN idno;  
    END;  
END;  
/  
/
```

Object types also have some important differences from the more familiar datatypes that are native to a relational database:

- A set of object types does not come ready-made with the database. Instead, you define the specific object types you want by extending built-in types with user-defined ones as shown in [Example 1-1](#).
- Object types are composed of parts, called attributes and methods.
  - Attributes hold the data about an object's features of interest. For example, a `student` object type might have `name`, `major`, and `graduation date` attributes. An attribute has a declared datatype which can in turn be another object type. Taken together, the attributes of an object instance contain that object's data.
  - Methods are procedures or functions provided to enable applications to perform useful operations on the attributes of the object type. Methods are an optional element of an object type. They define the behavior of objects of that type and determine what (if anything) that type of object can do.
  - Object types are less generic than native datatypes. In fact, this is one of their major virtues. You can define object types to model the actual structure of the real-world entities, such as customers and purchase orders, that application programs deal with. This can make it easier and more intuitive to manage the data for these entities. In this respect object types are like Java and C++ classes.

You can think of an object type as a blueprint or template which defines structure and behavior. An instantiation of the object type creates an object built according to the template. Object types are database schema objects, subject to the same kinds of administrative control as other schema objects.

**See Also:** [Chapter 6, "Managing Oracle Objects"](#)

You can use object types to model the actual structure of real-world objects. Object types enable you to capture the structural interrelationships of objects and their attributes instead of flattening this structure into a two-dimensional, purely relational schema of tables and columns. With object types you can store related pieces of data in a unit along with the behaviors defined for that data. Application code can then retrieve and manipulate these units as objects.

## Objects

When you create a variable of an object type, you create an instance of the type and the result is an object. An object has the attributes and methods defined for its type. Because an object instance is a concrete thing, you can assign values to its attributes and call its methods.

You use the `CREATE TYPE` statement to define object types. In [Example 1-1](#), the `CREATE TYPE` statement defines the object type `person_typ`.

The indented elements `idno`, `name`, and `phone` in the `CREATE TYPE` statements are attributes. Each has a datatype declared for it. These are simplified examples and do not show how to specify member methods.

Defining an object type does not allocate any storage. After they are defined, object types can be used in SQL statements in most of the same places you can use types like `NUMBER` or `VARCHAR2`.

For example, you might define a relational table to keep track of your contacts:

```
CREATE TABLE contacts (  
    contact          person_typ,  
    contact_date     DATE );  
  
INSERT INTO contacts VALUES (  
    person_typ (65, 'Vrinda Mills', '1-800-555-4412'), '24 Jun 2003' );
```

The `contacts` table is a relational table with an object type as the datatype of one of its columns. Objects that occupy columns of relational tables are called column objects.

**See Also:** ["Row Objects and Column Objects"](#) on page 1-9

## Object Methods

Methods are functions or procedures that you can declare in an object type definition to implement behavior that you want objects of that type to perform. For example, a method is declared in [Example 1-1](#) to allow comparisons between `person_typ` objects.

The general kinds of methods that can be declared in a type definition are:

- Member
- Static
- Constructor

A principal use of methods is to provide access to the data of an object. You can define methods for operations that an application is likely to want to perform on the data so that the application does not have to code these operations itself. To perform the operation, an application calls the appropriate method on the appropriate object.

For example, the following SQL statement uses the `get_idno()` method to display the Id number of persons in the `contacts` table:

```
SELECT c.contact.get_idno() FROM contacts c;
```

You can also define static methods to compare object instances and to perform operations that do not use any particular object's data but instead are global to an object type.

A constructor method is implicitly defined for every object type, unless this default constructor is over-written with a user-defined constructor. A constructor method is called on a type to construct or create an object instance of the type.

**See Also:** ["Object Methods"](#) on page 2-8

### Type Inheritance

Type inheritance adds to the usefulness of objects by enabling you to create type hierarchies by defining successive levels of increasingly specialized subtypes that derive from a common ancestor object type, which is called a supertype of the derived types. Derived subtypes inherit the features of the parent object type but extend the parent type definition. The specialized types can add new attributes or methods, or redefine methods inherited from the parent. The resulting type hierarchy provides a higher level of abstraction for managing the complexity of an application model.

For example, specialized types of persons, such as a student type or a part-time student type with additional attributes or methods, might be derived from a general person object type. See ["Inheritance in SQL Object Types"](#) on page 2-15.

### Type Evolution

Using an `ALTER TYPE` statement, you can modify, or evolve, an existing user-defined type to make the following changes:

- Add and drop attributes
- Add and drop methods
- Modify a numeric attribute to increase its length, precision, or scale
- Modify a varying length character attribute to increase its length
- Change a type's `FINAL` and `INSTANTIABLE` properties

Dependencies of a type to be altered are checked using essentially the same validations applied for a `CREATE TYPE` statement. If a type or any of its dependent types fails the type validations, the `ALTER TYPE` statement rolls back.

Metadata for all tables and columns that use an altered type are updated for the new type definition so that data can be stored in them in the new format. Existing data can be converted to the new format either all at once or piecemeal, as it is updated. In either case, data is always presented in the new type definition even if it is still stored in the format of the older one.

## Object Tables

An object table is a special kind of table in which each row represents an object. For example, the following statement creates an object table for `person_typ` objects:

```
CREATE TABLE person_obj_table OF person_typ;
```

You can view this table in two ways:

- As a single-column table in which each row is a `person_typ` object, allowing you to perform object-oriented operations
- As a multi-column table in which each attribute of the object type `person_typ`; such as `idno`, `name`, and `phone`; occupies a column, allowing you to perform relational operations

For example, you can execute the following instructions:

```
INSERT INTO person_obj_table VALUES (  
    1, 'John Smith', '1-800-555-1212');
```

```
SELECT VALUE(p) FROM person_obj_table p  
    WHERE p.name = 'John Smith';
```

The first statement inserts a `person_typ` object into `person_obj_table`, treating `person_table` as a multi-column table. The second selects from `person_obj_table` as a single-column table, using the `VALUE` function to return rows as object instances. See "[VALUE](#)" on page 2-41 for information on the `VALUE` function

By default, every row object in an object table has an associated logical object identifier (OID) that uniquely identifies it in an object table. In a distributed and replicated environment, the system-generated unique identifier lets Oracle identify objects unambiguously.

**See Also:** ["Storage Considerations for Object Identifiers \(OIDs\)"](#) on page 8-6 for information on Object Identifiers and using `REFs` to `OIDs`

### Row Objects and Column Objects

Objects that are stored in complete rows in object tables are called row objects. Objects that are stored as columns of a table in a larger row, or are attributes of other objects, are called column objects.

### Object Views

An object view is a way to access relational data using object-relational features. It lets you develop object-oriented applications without changing the underlying relational schema.

Oracle allows the creation of an object abstraction over existing relational data through the object view mechanism. You access objects that belong to an object view in the same way that you access row objects in an object table. Oracle also supports materialized view objects of user-defined types from data stored in relational schemas and tables. By using object views, you can develop object-oriented applications without having to modify existing relational database schemas.

Object views also let you exploit the polymorphism that a type hierarchy makes possible. A polymorphic expression can take a value of the expression's declared type or any of that type's subtypes. If you construct a hierarchy of object views that mirrors some or all of the structure of a type hierarchy, you can query any view in the hierarchy to access data at just the level of specialization you are interested in. If you query an object view that has subviews, you can get back polymorphic data—rows for both the type of the view and for its subtypes.

**See Also:** [Chapter 5, "Applying an Object Model to Relational Data"](#)

### References

A `REF` is a logical pointer to a row object that is constructed from the object identifier (OID) of the referenced object and is an Oracle built-in datatype. `REFs` and collections of `REFs` model associations among objects, particularly many-to-one relationships, thus reducing the need for foreign keys. `REFs` provide an easy mechanism for navigating between objects. You can use the dot notation to follow the pointers. Oracle does joins for you when needed, and in some cases can avoid doing joins.

You can use a REF to examine or update the object it refers to. You can also use a REF to obtain the object it refers to. You can change a REF so that it points to a different object of the same object type hierarchy or assign it a null value.

The following example illustrates a simple use of a REF.

**Example 1–2 Using a REF to an Object**

```
CREATE TYPE emp_person_typ AS OBJECT (  
    name    VARCHAR2(30),  
    manager REF emp_person_typ );  
/  
CREATE TABLE emp_person_obj_table OF emp_person_typ;  
  
INSERT INTO emp_person_obj_table VALUES (  
    emp_person_typ ('John Smith', NULL));  
  
INSERT INTO emp_person_obj_table  
    SELECT emp_person_typ ('Bob Jones', REF(e))  
    FROM emp_person_obj_table e  
    WHERE e.name = 'John Smith';
```

**See Also:**

- ["Rules for REF Columns and Attributes"](#) on page 2-6
- ["Design Considerations for REFs"](#) on page 8-8

**Scoped REFs**

In declaring a column type, collection element, or object type attribute to be a REF, you can constrain it to contain only references to a specified object table. Such a REF is called a *scoped REF*. Scoped REF types require less storage space and allow more efficient access than unscoped REF types.

The following example shows REF column `contact_ref` scoped to `person_obj_table` which is an object table of type `person_typ`.

```
CREATE TABLE contacts_ref (  
    contact_ref  REF person_typ SCOPE IS person_obj_table,  
    contact_date DATE );
```

To insert a row in the table, you could issue the following:

```
INSERT INTO contacts_ref  
    SELECT REF(p), '26 Jun 2003'
```

```
FROM person_obj_table p
WHERE p.idno = 1;
```

A REF can be scoped to an object table of the declared type (`person_typ` in the example) or of any subtype of the declared type. If scoped to an object table of a subtype, the REF column is effectively constrained to hold references only to instances of the subtype (and its subtypes, if any) in the table.

**See Also:** ["Inheritance in SQL Object Types"](#) on page 2-15

### Dangling REFs

It is possible for the object identified by a REF to become unavailable through either deletion of the object or a revoking of privileges. Such a REF is called dangling. Oracle SQL provides a predicate (called `IS DANGLING`) to allow testing REFs for this condition.

Dangling REFs can be avoided by defining referential integrity constraints. See ["Rules for REF Columns and Attributes"](#) on page 2-6.

### Dereferencing REFs

Accessing the object referred to by a REF is called dereferencing the REF. Oracle provides the `DEREF` operator to do this. Dereferencing a dangling REF returns a null object. For example:

```
SELECT DEREF(c.contact_ref), c.contact_date FROM contacts_ref c;
```

Oracle also provides implicit dereferencing of REFs. For example, to access the manager's name for an employee, you can use a SQL expression similar to the following:

```
SELECT e.name, e.manager.name FROM emp_person_obj_table
WHERE e.name = 'Bob Jones';
```

In the example, `e.manager.name` follows the pointer from the person's manager, and retrieves the manager's name. Following the REF like this is allowed in SQL, but not in PL/SQL.

### Obtaining REFs

You can obtain a REF to a row object by selecting the object from its object table and applying the REF operator. For example, you can obtain a REF to the person with identification number 1 as follows:

```
DECLARE
```

```
    person_ref REF person_typ;
BEGIN
    SELECT REF(p) INTO person_ref
        FROM person_obj_table p
        WHERE p.idno = 1;
END;
/
```

The query must return exactly one row.

**See Also:** ["Storage Size of REFs"](#) on page 8-8

### Collections

For modeling multi-valued attributes and many to many relationships, Oracle supports two collection datatypes: varrays and nested tables. Collection types can be used anywhere other datatypes can be used. You can have object attributes of a collection type in addition to columns of a collection type. For example, you might give a purchase order object type a nested table attribute to hold the collection of line items for a given purchase order.

You use the `CREATE TYPE` statement to define collection types. In [Example 1-3](#), the `CREATE TYPE` statements define the object types `people_typ` and `dept_persons_typ`.

#### **Example 1-3** *Creating a Collection Datatype*

```
CREATE TYPE people_typ AS TABLE OF person_typ;
/

CREATE TYPE dept_persons_typ AS OBJECT (
    dept_no    CHAR(5),
    dept_name  CHAR(20),
    dept_mgr   person_typ,
    dept_emps  people_typ);
/
```

In this simplified example, `people_typ` is a collection type, specifically a nested table type. The `dept_persons_typ` object type has an attribute `people_typ` of this type. Each row in the `people_typ` nested table is an object of type `person_typ` which was defined in [Example 1-1](#) on page 1-4.

**See Also:** ["Creating Collection Datatypes"](#) on page 3-2

## Language Binding Features

This section lists the key features of the object-relational model that are related to languages and application programming interfaces (APIs).

### SQL Object Extensions

To support the new object-related features, SQL extensions, including new DDL, have been added to create, alter, or drop object types; to store object types in tables; and to create, alter, or drop object views. There are DML and query extensions to support object types, references, and collections.

**See Also:** ["SQL"](#) on page 4-2

### PL/SQL Object Extensions

PL/SQL is an Oracle database programming language that is tightly integrated with SQL. With the addition of user-defined types and other SQL types, PL/SQL has been enhanced to operate on user-defined types seamlessly. Thus, application developers can use PL/SQL to implement logic and operations on user-defined types that execute in the database server.

**See Also:** ["PL/SQL"](#) on page 4-2

### Java Support for Oracle Objects

Oracle Java VM is tightly integrated with the RDBMS and supports access to Oracle Objects through object extensions to Java Database Connectivity (JDBC), which provides dynamic SQL, and SQLJ, which provides static SQL. Thus, application developers can use the Java to implement logic and operations on user-defined types that execute in the database server. With Oracle, you can now also create SQL types mapped to existing Java classes to provide persistent storage for Java objects using SQLJ object types where all the methods are implemented in their corresponding Java classes. See ["Java Object Storage"](#) on page 4-16.

**See Also:** *Oracle Database JDBC Developer's Guide and Reference*

### External Procedures

Database functions, procedures, or member methods of an object type can be implemented in PL/SQL, Java, or C as external procedures. External procedures are best suited for tasks that are more quickly or easily done in a low-level language such as C, which is more efficient at machine-precision calculation. External procedures are always run in a safe mode outside the address space of the RDBMS

server. Generic external procedures can be written that declare one or more parameters to be of a system-defined generic type. The generic type permits a procedure that uses it to work with data of any built-in or user-defined type.

### **Object Type Translator/JPublisher**

Object Type Translator (OTT) and Oracle JPublisher provide client-side mappings to object type schemas by using schema information from the Oracle data dictionary to generate header files containing Java classes and C structures and indicators. These generated header files can be used in host-language applications for transparent access to database objects.

### **Client-Side Cache**

Oracle provides an object cache for efficient access to persistent objects stored in the database. Copies of objects can be brought into the object cache. Once the data has been cached in the client, the application can traverse through these at memory speed. Any changes made to objects in the cache can be committed to the database by using the object extensions to Oracle® Call Interface programmatic interfaces.

### **Oracle Call Interface and Oracle C++ Call Interface**

Oracle Call Interface (OCI) and Oracle C++ Call Interface provide a comprehensive application programming interface for application and tool developers seeking to use the object capabilities of Oracle. Oracle Call Interface provides a run-time environment with functions to connect to an Oracle server, and control transactions that access objects in the server. It allows application developers to access and manipulate objects and their attributes in the client-side object cache either navigationally, by traversing a graph of inter-connected objects, or associatively by specifying the nature of the data through declarative SQL DML. Oracle Call Interface also provides a number of functions for accessing metadata information at run-time about object types defined in the server. Such a set of functions facilitates dynamic access to the object metadata and the actual object data stored in the database.

**See Also:** ["Oracle Call Interface \(OCI\)"](#) on page 4-2 and ["Oracle C++ Call Interface \(OCCI\)"](#) on page 4-8

### **Pro\*C/C++ Object Extensions**

The Oracle Pro\*C™ precompiler provides an embedded SQL application programming interface and offers a higher level of abstraction than Oracle Call Interface. Like Oracle Call Interface, the Pro\*C precompiler allows application developers to use the Oracle client-side object cache and the Object Type Translator

Utility. Pro\*C supports the use of C bind variables for Oracle object types. Furthermore, Pro\*C provides new simplified syntax to allocate and free objects of SQL types and access them by either SQL DML, or through the navigational interface. Thus, it provides application developers many benefits, including compile-time type checking of (client-side) bind variables against the schema in the server, automatic mapping of object data in an Oracle server to program bind variables in the client, and simple ways to manage and manipulate database objects in the client process.

**See Also:** ["Oracle Call Interface \(OCI\)"](#) on page 4-2

### **OO4O Object Extensions**

Oracle Objects For OLE (OO4O) is a set of COM Automation interfaces/objects for connecting to Oracle database servers, executing queries and managing the results. Automation interfaces in OO4O provide easy and efficient access to Oracle features and can be used from virtually any programming or scripting language that supports the Microsoft COM Automation technology. This includes Visual Basic, Visual C++, VBA in Excel, VBScript and JavaScript in IIS Active Server Pages.

**See Also:** ["Oracle Objects For OLE \(OO4O\)"](#) on page 4-10



---

# Basic Components of Oracle Objects

This chapter provides basic information about working with objects. It explains what object types, methods, and collections are and describes how to create and work with a hierarchy of object types that are derived from a shared root type and are connected by inheritance.

This chapter contains these topics:

- [SQL Object Types and References](#)
- [Object Methods](#)
- [Inheritance in SQL Object Types](#)
- [Functions and Operators Useful with Objects](#)

## SQL Object Types and References

This section describes SQL object types and references, including:

- [Null Objects and Attributes](#)
- [Character Length Semantics](#)
- [Constraints for Object Tables](#)
- [Indexes for Object Tables](#)
- [Triggers for Object Tables](#)
- [Rules for REF Columns and Attributes](#)
- [Name Resolution](#)

You can create a SQL object type with the `CREATE TYPE` statement. An example of creating an object type is shown in [Example 1-1](#) on page 1-4.

### Null Objects and Attributes

A table column, object attribute, collection, or collection element is `NULL` if it has been initialized to `NULL` or has not been initialized at all. Usually, a `NULL` value is replaced by an actual value later on.

An object whose value is `NULL` is called atomically null. An atomically null object is different from one that simply happens to have null values for all its attributes. When all the attributes of an object are null, these attributes can still be changed, and the object's methods can be called. With an atomically null object, you can do neither of these things.

For example, consider the `contacts` table which contains the `person_typ` object type defined in [Example 1-1](#) on page 1-4.

```
CREATE TABLE contacts (  
  contact      person_typ,  
  contact_date DATE );
```

#### The statement

```
INSERT INTO contacts VALUES (  
  person_typ (NULL, NULL, NULL),  
  '24 Jun 2003' );
```

gives a different result from

```
INSERT INTO contacts VALUES (  
  person_typ (NULL, NULL, NULL),  
  '24 Jun 2003' );
```

```
NULL,  
'24 Jun 2003' );
```

In both cases, Oracle allocates space in `contacts` for a new row and sets its `DATE` column to the value given. But in the first case, Oracle allocates space for an object in the `contact` column and sets each of the object's attributes to `NULL`. In the second case, Oracle sets the `person_typ` field itself to `NULL` and does not allocate space for an object.

In some cases, you can omit checks for null values. A table row or row object cannot be null. A nested table of objects cannot contain an element whose value is `NULL`.

A nested table or array can be null, so you do need to handle that condition. A null collection is different from an empty one, that is, a collection containing no elements.

## Character Length Semantics

Lengths for character types `CHAR` and `VARCHAR2` may be specified as a number of characters, instead of bytes, in object attributes and collections even if some of the characters consist of multiple bytes.

To specify character-denominated lengths for `CHAR` and `VARCHAR2` attributes, you add a qualifier `char` to the length specification.

Like `CHAR` and `VARCHAR2`, `NCHAR` and `NVARCHAR2` may also be used as attribute types in objects and collections. These types are always implicitly measured in terms of characters, so no `char` qualifier is used.

For example, the following statement creates an object with both a character-length `VARCHAR2` attribute and an `NCHAR` attribute:

```
CREATE TYPE employee_typ AS OBJECT (  
    name          VARCHAR2(30 char),  
    language      NCHAR(10),  
    phone         VARCHAR2(20) );  
/
```

For `CHAR` and `VARCHAR2` attributes whose length is specified without a `char` qualifier (for example, the `phone` attribute), the default unit of measure characters or bytes is determined by whether the `NLS_LENGTH_SEMANTICS` initialization parameter is set to `CHAR` or `BYTE`.

**See Also:** *Oracle Database Globalization Support Guide* for information on character length semantics

## Constraints for Object Tables

You can define constraints on an object table just as you can on other tables. You can define constraints on the leaf-level scalar attributes of a column object, with the exception of `REFs` that are not scoped.

The following examples illustrate the possibilities. The first example places a `PRIMARY KEY` constraint on the `idno` column of the object table `person_extent`:

```
CREATE TYPE location_typ (  
    building_no NUMBER,  
    city          VARCHAR2(40) );  
/  
  
CREATE TYPE office_typ AS OBJECT (  
    office_id     VARCHAR(10),  
    office_loc    location_typ,  
    occupant      person_typ );  
/  
  
CREATE TABLE office_tab OF office_typ (  
    office_id     PRIMARY KEY );
```

The `department_mgrs` table in the next example has a column whose type is the object type `location_typ` defined in the previous example. The example defines constraints on scalar attributes of the `location_typ` objects that appear in the `dept_loc` column of the table.

```
CREATE TABLE department_mgrs (  
    dept_no      NUMBER PRIMARY KEY,  
    dept_name    CHAR(20),  
    dept_mgr     person_typ,  
    dept_loc     location_typ,  
    CONSTRAINT  dept_loc_cons1  
        UNIQUE (dept_loc.building_no, dept_loc.city),  
    CONSTRAINT  dept_loc_cons2  
        CHECK (dept_loc.city IS NOT NULL) );  
  
INSERT INTO department_mgrs VALUES  
    ( 101, 'Physical Sciences',  
      person_typ(65, 'Vrinda Mills', '1-800-555-4412'),  
      location_typ(300, 'Palo Alto'));
```

## Indexes for Object Tables

You can define indexes on an object table or on the storage table for a nested table column or attribute just as you can on other tables.

You can define indexes on leaf-level scalar attributes of column objects, as shown in the following example. You can only define indexes on REF attributes or columns if the REF is scoped.

Here, `dept_addr` is a column object, and `city` is a leaf-level scalar attribute of `dept_addr` that we want to index:

```
CREATE TABLE department_loc (
  dept_no    NUMBER PRIMARY KEY,
  dept_name  CHAR(20),
  dept_addr  location_typ );

CREATE INDEX i_dept_addr1
  ON department_loc (dept_addr.city);
```

Wherever Oracle expects a column name in an index definition, you can also specify a scalar attribute of an object column.

## Triggers for Object Tables

You can define triggers on an object table just as you can on other tables. You cannot define a trigger on the storage table for a nested table column or attribute. You cannot modify LOB values in a trigger body. Otherwise, there are no special restrictions on using object types with triggers.

The following example defines a trigger on the `office_tab` table defined in ["Constraints for Object Tables"](#) on page 2-4.

```
CREATE TABLE movement (
  idno        NUMBER,
  old_office  location_typ,
  new_office  location_typ );

CREATE TRIGGER trigger1
  BEFORE UPDATE
    OF office_loc
    ON office_tab
  FOR EACH ROW
  WHEN (new.office_loc.city = 'Redwood Shores')
  BEGIN
  IF :new.office_loc.building_no = 600 THEN
```

```
        INSERT INTO movement (idno, old_office, new_office)
        VALUES (:old.occupant.idno, :old.office_loc, :new.office_loc);
    END IF;
END;
```

## Rules for REF Columns and Attributes

In Oracle, a REF column or attribute can be unconstrained or constrained using a SCOPE clause or a referential constraint clause. When a REF column is unconstrained, it may store object references to row objects contained in any object table of the corresponding object type.

Oracle does not ensure that the object references stored in such columns point to valid and existing row objects. Therefore, REF columns may contain object references that do not point to any existing row object. Such REF values are referred to as dangling references.

A REF column may be constrained to be scoped to a specific object table. All the REF values stored in a column with a SCOPE constraint point at row objects of the table specified in the SCOPE clause. The REF values may, however, be dangling.

A REF column may be constrained with a REFERENTIAL constraint similar to the specification for foreign keys. The rules for referential constraints apply to such columns. That is, the object reference stored in these columns must point to a valid and existing row object in the specified object table.

PRIMARY KEY constraints cannot be specified for REF columns. However, you can specify NOT NULL constraints for such columns.

## Name Resolution

Oracle SQL lets you omit qualifying table names in some relational operations. For example, if dept\_addr is a column in the department\_loc table and old\_office is a column in the movement table, you can write:

```
SELECT * FROM department_loc WHERE EXISTS
    (SELECT * FROM movement WHERE dept_addr = old_office);
```

Oracle determines which table each column belongs to.

Using the dot notation, you can qualify the column names with table names or table aliases to make things more maintainable. For example:

```
SELECT * FROM department_loc WHERE EXISTS
    (SELECT * FROM movement WHERE department_loc.dept_addr = movement.old_office);
```

```
SELECT * FROM department_loc d WHERE EXISTS
  (SELECT * FROM movement m WHERE d.dept_addr = m.old_office);
```

In some cases, object-relational features require you to specify the table aliases.

### When Table Aliases Are Required

Using unqualified names can lead to problems. If you add an `assignment` column to `depts` and forget to change the query, Oracle automatically recompiles the query such that the inner `SELECT` uses the `assignment` column from the `depts` table. This situation is called inner capture.

To avoid inner capture and similar problems resolving references, Oracle requires you to use a table alias to qualify any dot-notational reference to methods or attributes of objects. Use of a table alias is optional when referencing top-level attributes of an object table directly, without using the dot notation.

For example, the following statements define two tables that contain the `person_typ` object type. `person_obj_table` is an object table for objects of type `person_typ`, and `contacts` is a relational table that contains a column of an object type.

The following queries show some correct and incorrect ways to reference attribute `idno`:

```
SELECT idno FROM person_obj_table;           --Correct
SELECT contact.idno FROM contacts;          --Illegal
SELECT contacts.contact.idno FROM contacts; --Illegal
SELECT p.contact.idno FROM contacts p;      --Correct
```

- In the first `SELECT` statement, `idno` is the name of a column of `person_obj_table`. It references this top-level attribute directly, without using the dot notation, so no table alias is required.
- In the second `SELECT` statement, `idno` is the name of an attribute of the `person_typ` object in the column named `contact`. This reference uses the dot notation and so requires a table alias, as shown in the fourth `SELECT` statement.
- The third `SELECT` uses the table name itself to qualify this the reference. This is incorrect; a table alias is required.

You must qualify a reference to an object attribute or method with a table alias rather than a table name even if the table name is itself qualified by a schema name.

For example, the following expression tries to refer to the `objtest` schema, `department_loc` table, `dept_addr` column, and `city` attribute of that column.

But the expression is incorrect because `department_loc` is a table name, not an alias.

```
objtest.department_loc.dept_addr.city
```

The same requirement applies to attribute references that use `REFs`.

Table aliases should uniquely pick out the same table throughout a query and should not be the same as schema names that could legally appear in the query.

---

---

**Note:** Oracle recommends that you define table aliases in all `UPDATE`, `DELETE`, and `SELECT` statements and subqueries and use them to qualify column references whether or not the columns contain object types.

---

---

## Restriction on Using User-Defined Types with a Remote Database

User-defined types (specifically, types declared with a `SQL CREATE TYPE` statement, as opposed to types declared within a `PL/SQL` package) are currently useful only within a single database. You cannot use a database link to do any of the following:

- Connect to a remote database to select, insert, or update a user-defined type or an object `REF` on a remote table

You can use the `CREATE TYPE` statement with the optional keyword `OID` to create a user-specified object identifier (OID) that allows an object type to be used in multiple databases. See the discussion on assigning an OID to an object type in the *Oracle Data Cartridge Developer's Guide*.

- Use database links within `PL/SQL` code to declare a local variable of a remote user-defined type
- Convey a user-defined type argument or return value in a `PL/SQL` remote procedure call.

## Object Methods

Methods are functions or procedures that you can declare in an object type definition to implement behavior that you want objects of that type to perform. An application calls the methods to invoke the behavior.

For example, you might declare a method `get_sum()` to get a purchase order object to return the total cost of its line items. The following line of code calls such a method for purchase order `po` and returns the amount into `sum_line_items`:

```
sum_line_items = po.get_sum();
```

In SQL, the parentheses are required for all method calls. Unlike with PL/SQL functions and procedures, SQL requires parentheses for method calls that do not have arguments.

Methods can be written in PL/SQL or virtually any other programming language. Methods written in PL/SQL or Java are stored in the database. Methods written in other languages, such as C, are stored externally.

The topics described in this section are:

- [Member Methods](#)
- [Methods for Comparing Objects](#)
- [Static Methods](#)
- [Constructor Methods](#)
- [External Implemented Methods](#)

## Member Methods

Member methods are the means by which an application gains access to an object instance's data. You define a member method in the object type for each operation that you want an object of that type to be able to perform. For example, the method `get_sum()` that sums the total cost of a purchase order's line items operates on the data of a particular purchase order and is a member method.

Member methods have a built-in parameter named `SELF` that denotes the object instance on which the method is currently being invoked. Member methods can reference the attributes and methods of `SELF` without a qualifier. This makes it simpler to write member methods. For example, the following code shows a method declaration that takes advantage of `SELF` to omit qualification of the attributes `num` and `den`:

```
CREATE TYPE rational_typ AS OBJECT (  
    num INTEGER,  
    den INTEGER,  
    MEMBER PROCEDURE normalize_proc);  
/
```

```
CREATE TYPE BODY rational_typ AS
  MEMBER PROCEDURE normalize_proc IS
    g INTEGER;
  BEGIN
    g := gcd(SELf.num, SELf.den);
    g := gcd(num, den);          -- equivalent to previous line
    num := num / g;
    den := den / g;
  END normalize_proc;
END;
/
```

`SELF` does not need to be explicitly declared, although it can be. It is always the first parameter passed to the method.

- In member functions, if `SELF` is not declared, its parameter mode defaults to `IN`.
- In member procedures, if `SELF` is not declared, its parameter mode defaults to `IN OUT`. The default behavior does not include the `NOCOPY` compiler hint.

You can invoke a member method using the dot notation `object_variable.method()`. This notation specifies the object on which to invoke the method, then the method to call. Any parameters must be placed inside the required parentheses.

**See Also:** ["Using SELF IN OUT NOCOPY with Member Procedures"](#) on page 8-28

## Methods for Comparing Objects

The values of a scalar datatype such as `CHAR` or `REAL` have a predefined order, which allows them to be compared. But an object type, such as a `customer_typ`, which can have multiple attributes of various datatypes, has no predefined axis of comparison. To be able to compare and order variables of an object type, you must specify a basis for comparing them.

Two special kinds of member methods can be defined for doing this: map methods and order methods.

### Map Methods

A map method is an optional kind of method that provides a basis for comparing objects by mapping object instances to one of the scalar types `DATE`, `NUMBER`, `VARCHAR2` or to an ANSI SQL type such as `CHARACTER` or `REAL`. With a map

method, you can order any number of objects by calling each object's map method once to map that object to a position on the axis used for the comparison, such as a number or date. [Example 1-1](#) on page 1-4 contains a simple map method.

From the standpoint of writing one, a map method is simply a parameter-less member function that uses the `MAP` keyword and returns one of the datatypes just listed. What makes a map method special is that, if an object type defines one, the method is called automatically to evaluate such comparisons as `obj_1 > obj_2` and comparisons implied by the `DISTINCT`, `GROUP BY`, `UNION`, and `ORDER BY` clauses which require sorting by rows.

Where `obj_1` and `obj_2` are two object variables that can be compared using a map method `map()`, the comparison:

```
obj_1 > obj_2
```

is equivalent to:

```
obj_1.map() > obj_2.map()
```

And similarly for other relational operators besides the greater than (`>`) operator.

The following example defines a map method `area()` that provides a basis for comparing rectangle objects by their area:

```
CREATE TYPE rectangle_typ AS OBJECT (
    len NUMBER,
    wid NUMBER,
    MAP MEMBER FUNCTION area RETURN NUMBER);
/

CREATE TYPE BODY rectangle_typ AS
    MAP MEMBER FUNCTION area RETURN NUMBER IS
    BEGIN
        RETURN len * wid;
    END area;
END;
/
```

An object type can declare at most one map method or one order method. A subtype can declare a map method only if its root supertype declares one.

**See Also:** ["Equal and Not Equal Comparisons"](#) on page 3-17 for the use of map methods when comparing collections that contain object types

## Order Methods

Order methods make direct object-to-object comparisons. Unlike map methods, they cannot map any number of objects to an external axis. They simply tell you that the current object is less than, equal to, or greater than the other object that it is being compared to, with respect to the criterion used by the method.

An order method is a function with one declared parameter for another object of the same type. The method must be written to return either a negative number, zero, or a positive number. The return signifies that the object picked out by the `SELF` parameter is respectively less than, equal to, or greater than the other parameter's object.

As with map methods, an order method, if one is defined, is called automatically whenever two objects of that type need to be compared.

Order methods are useful where comparison semantics may be too complex to use a map method. For example, to compare binary objects such as images, you might create an order method to compare the images by their brightness or number of pixels.

An object type can declare at most one order method or one map method. Only a type that is not derived from another type can declare an order method; a subtype cannot define one.

The following example shows an order method that compares locations by building number:

```
CREATE TYPE location_typ AS OBJECT (  
    building_no NUMBER,  
    city          VARCHAR2(40),  
    ORDER MEMBER FUNCTION match (l location_typ) RETURN INTEGER );  
/  
CREATE TYPE BODY location_typ AS  
    ORDER MEMBER FUNCTION match (l location_typ) RETURN INTEGER IS  
    BEGIN  
        IF building_no < l.building_no THEN  
            RETURN -1;                -- any negative number will do  
        ELSIF building_no > l.building_no THEN  
            RETURN 1;                -- any positive number will do  
        ELSE  
            RETURN 0;  
        END IF;  
    END;  
END;  
/
```

## Guidelines for Comparison Methods

A map method maps object values into scalar values and can order multiple values by their position on the scalar axis. An order method directly compares values for two particular objects.

You can declare a map method or an order method but not both. If you declare a method of either type, you can compare objects in SQL and procedural statements. However, if you declare neither method, you can compare objects only in SQL statements and only for equality or inequality. (Two objects of the same type count as equal only if the values of their corresponding attributes are equal.)

When sorting or merging a large number of objects, use a map method. One call maps all the objects into scalars, then sorts the scalars. An order method is less efficient because it must be called repeatedly (it can compare only two objects at a time).

**See Also:** ["Performance of Object Comparisons"](#) on page 8-7

## Comparison Methods in Type Hierarchies

In a type hierarchy, where definitions of specialized types are derived from definitions of more general types, only the root type—the most basic type, from which all other types are derived—can define an order method. If the root type does not define one, its subtypes cannot define one either.

If the root type specifies a map method, any of its subtypes can define a map method that overrides the map method of the root type. But if the root type does not specify a map method, no subtype can specify one either.

So if the root type does not specify either a map or an order method, none of the subtypes can specify either a map or order method.

**See Also:**

- ["Inheritance in SQL Object Types"](#) on page 2-15
- ["Inheriting, Overloading, and Overriding Methods"](#) on page 2-20

## Static Methods

Static methods are invoked on the object type, not its instances. You use a static method for operations that are global to the type and do not need to reference the data of a particular object instance. A static method has no `SELF` parameter.

You invoke a static method by using the dot notation to qualify the method call with the name of the object type, such as:

```
type_name.method()
```

**See Also:** ["Static Methods"](#) on page 8-27 for information on design considerations

## Constructor Methods

Every object type has a constructor method implicitly defined for it by the system. A constructor method is a function that returns a new instance of the user-defined type and sets up the values of its attributes.

The system implicitly defines a constructor function called the attribute value constructor for all object types that have attributes.

You can also define constructor functions of your own called user-defined constructors to create and initialize objects of such types. Attribute value constructors are convenient to use because they already exists, but user-defined constructors have some important advantages with respect to type evolution.

**See Also:** ["Advantages of User-Defined Constructors"](#) on page 7-20 for information on user-defined constructors and their advantages

The name of the constructor method is just the name of the object type, as shown in the following simple example:

```
person_typ (1, 'John Smith', '1-800-555-1212'),
```

A literal invocation of a constructor method is a call to the constructor method in which any arguments are either literals, or further literal invocations of constructor methods.

For example, consider the `people_typ` object type that is a table of `person_typ` that is defined in [Example 1-1](#) on page 1-4.

```
CREATE TYPE people_typ AS TABLE OF person_typ;  
/
```

**See Also:** ["Constructor Methods for Collections"](#) on page 3-2 for information on user-defined constructors and their advantages

## External Implemented Methods

You can use PL/SQL to invoke external subprograms that have been written in other languages. This provides access to the strengths and capabilities of those languages.

**See Also:** *PL/SQL User's Guide and Reference* for information on external implemented methods

## Inheritance in SQL Object Types

Object types enable you to model the real-world entities such as customers and purchase orders that your application works with. But this is just the first step in exploiting the capabilities of objects. With objects, you cannot only model an entity such as a customer, you can also define different specialized types of customers in a type hierarchy under the original type. You can then perform operations on a hierarchy and have each type implement and execute the operation in a special way.

A type hierarchy is a sort of family tree of object types. It consists of a parent base type, called a supertype, and one or more levels of child object types, called subtypes, derived from the parent.

Subtypes in a hierarchy are connected to their supertypes by inheritance. This means that subtypes automatically acquire the attributes and methods of their parent type. It also means that subtypes automatically acquire any changes made to these attributes or methods in the parent: any attributes or methods updated in a supertype are updated in subtypes as well.

A subtype becomes a specialized version of the parent type by adding new attributes and methods to the set inherited from the parent or by redefining methods it inherits. Redefining an inherited methods gives a subtype its own way of executing the method. Add to this that an object instance of a subtype can generally be substituted for an object instance of any of its supertypes in code, and you have polymorphism.

Polymorphism is the ability of a slot for a value in code to contain a value of either a certain declared type or any of a range of the declared type's subtypes. A method called on whatever value occupies the slot may execute differently depending on the value's type because the various types might implement the method differently.

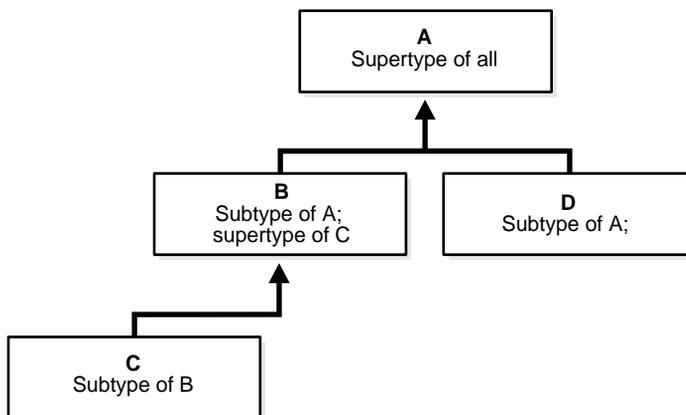
## Types and Subtypes

A subtype can be derived from a supertype either directly, or indirectly through intervening levels of other subtypes. A subtype can directly derive only from a

single supertype: it cannot derive jointly from more than one. A supertype can have multiple sibling subtypes, but a subtype can have at most one direct parent supertype. In other words, Oracle supports only single inheritance, not multiple inheritance.

A subtype is derived from a supertype by defining a specialized variant of the supertype. For example, from a `person_typ` object type you might derive the specialized types `student_typ` and `employee_typ`. Each of these subtypes is still at bottom a `person_typ`, but a special kind of person. What makes a subtype special and distinguishes it from its parent supertype is some change made in the subtype to the attributes or methods that the subtype received from its parent.

**Figure 2–1 Subtypes**



An object type's attributes and methods make the type what it is: they are its essential, defining features. If a `person_typ` object type has the three attributes `idno`, `name`, and `phone` and the method `get_idno()`, then any object type that is derived from `person_typ` will have these same three attributes and a method `get_idno()`. A subtype is a special case of its parent type, not a totally different kind of thing. As such, it shares with its parent type the features that make the general type what it is.

You can specialize the attributes or methods of a subtype in these ways:

- Add new attributes that its parent supertype does not have.

For example, you might specialize `student_typ` as a special kind of `person_typ` by adding to its definition an attribute for `major`. A subtype cannot drop

or change the type of an attribute it inherited from its parent; it can only add new attributes.

- Add entirely new methods that the parent does not have.
- Change the implementation of some of the methods a subtype inherits from its parent so that the subtype's version executes different code from the parent's.

For example, a `shape` object type might define a method `calculate_area()`. Two subtypes of `shape`, `rectilinear_shape`, and `circular_shape`, might each implement this method in a different way. See "[Inheriting, Overloading, and Overriding Methods](#)" on page 2-20.

Attributes and methods that a subtype gets from its parent type are said to be inherited. This means more than just that the attributes and methods are patterned on the parent's when the subtype is defined. With object types, the inheritance link remains live. Any changes made later on to the parent type's attributes or methods are also inherited so that the changes are reflected in the subtype as well. Unless a subtype reimplements an inherited method, it always contains the same core set of attributes and methods that are in the parent type, plus any attributes and methods that it adds.

Remember, a child type is not a different type from its parent; it is a particular kind of that type. If the general definition of `person_typ` ever changes, the definition of `student_typ` changes also.

The inheritance relationship that holds between a supertype and its subtypes is the source of both much of the power of objects and much of their complexity. It is a very powerful feature to be able to change a method in a supertype and have the change take effect in all the subtypes downstream just by recompiling. But this same capability means that you have to think about such things as whether you want to allow a type to be specialized or a method to be redefined. Similarly, it is a powerful feature for a table or column to be able to contain any type in a hierarchy, but then you must decide whether to allow this in a particular case, and you may need to constrain DML statements and queries so that they pick out from the type hierarchy just the range of types that you want. The following sections address these aspects of working with objects.

## FINAL and NOT FINAL Types and Methods

The definition of an object type determines whether subtypes can be derived from that type. To permit subtypes, the object type must be defined as not final. This is done by including the `NOT FINAL` keyword in its type declaration. For example:

```
CREATE TYPE person_typ AS OBJECT (
```

```
        idno          NUMBER,  
        name          VARCHAR2(30),  
        phone         VARCHAR2(20))  
NOT FINAL;  
/
```

The preceding statement declares `person_typ` to be a not final type such that subtypes of `person_typ` can be defined. By default, an object type is declared as final and subtypes cannot be derived from it.

You can change a final type to a not final type and vice versa with an `ALTER TYPE` statement. For example, the following statement changes `person_typ` to a final type:

```
ALTER TYPE person_typ FINAL;
```

You can alter a type from `NOT FINAL` to `FINAL` only if the target type has no subtypes.

Methods, too, can be declared to be final or not final. If a method is declared to be final, subtypes cannot override it by providing their own implementation. Unlike types, methods are not final by default and must be explicitly declared to be final.

The following statement creates a not final type containing a final member function:

```
CREATE TYPE person_typ AS OBJECT (  
    idno          NUMBER,  
    name          VARCHAR2(30),  
    phone         VARCHAR2(20)  
    FINAL MAP MEMBER FUNCTION get_idno RETURN NUMBER)  
NOT FINAL;  
/
```

**See Also:** ["Redefining Methods"](#) on page 2-21

## Creating Subtypes

You create a subtype using a `CREATE TYPE` statement that specifies the immediate parent of the subtype with an `UNDER` parameter:

```
CREATE TYPE student_typ UNDER person_typ (  
    dept_id NUMBER,  
    major VARCHAR2(30))  
NOT FINAL;  
/
```

The preceding statement creates `student_typ` as a subtype of `person_typ`. As a subtype of `person_typ`, `student_typ` inherits all the attributes declared in or inherited by `person_typ` and any methods inherited by `person_typ` or declared in `person_typ`.

The statement that defines `student_typ` specializes `person_typ` by adding two new attributes. New attributes declared in a subtype must have names that are different from the names of any attributes or methods declared in any of its supertypes, higher up in its type hierarchy.

A type can have multiple child subtypes, and these can also have subtypes. The following statement creates another subtype `employee_typ` under `person_typ`.

```
CREATE TYPE employee_typ UNDER person_typ (
    emp_id NUMBER,
    mgr VARCHAR2(30));
/
```

A subtype can be defined under another subtype. Again, the new subtype inherits all the attributes and methods that its parent type has, both declared and inherited. For example, the following statement defines a new subtype `part_time_student_typ` under `student_typ`. The new subtype inherits all the attributes and methods of `student_typ` and adds another attribute.

```
CREATE TYPE part_time_student_typ UNDER student_typ (
    number_hours NUMBER);
/
```

## NOT INSTANTIABLE Types and Methods

A type can be declared to be `NOT INSTANTIABLE`. If a type is not instantiable, there is no constructor (default or user-defined) for it, and you cannot instantiate instances of that type (objects, in other words). You might use this option with types that you intend to use solely as supertypes of specialized subtypes that you do instantiate. For example:

```
CREATE TYPE address_typ AS OBJECT(...) NOT INSTANTIABLE NOT FINAL;
CREATE TYPE USaddress_typ UNDER address_typ(...);
CREATE TYPE Intladdress_typ UNDER address_typ(...);
```

A method can also be declared to be not instantiable. Use this option when you want to declare a method in a type without implementing the method there. A type that contains a non-instantiable method must itself be declared not instantiable. For example:

```
CREATE TYPE person_typ AS OBJECT (  
    idno          NUMBER,  
    name          VARCHAR2(30),  
    phone         VARCHAR2(20),  
    NOT INSTANTIABLE MEMBER FUNCTION get_idno() RETURN NUMBER )  
    NOT INSTANTIABLE NOT FINAL;  
/
```

A non-instantiable method serves as a placeholder. You might define a non-instantiable method when you expect every subtype to override the method in a different way. In such a case, there is no point in defining the method in the supertype.

If a subtype does not provide an implementation for every inherited non-instantiable method, the subtype itself, like the supertype, must be declared not instantiable. A non-instantiable subtype can be defined under an instantiable supertype.

You can alter an instantiable type to a non-instantiable type and vice versa with an `ALTER TYPE` statement. In the following example, the `ALTER TYPE` statement makes `person_typ` instantiable:

```
CREATE TYPE person_typ AS OBJECT (  
    idno          NUMBER,  
    name          VARCHAR2(30),  
    phone         VARCHAR2(20))  
    NOT INSTANTIABLE NOT FINAL;  
/  
ALTER TYPE person_typ INSTANTIABLE;
```

You can alter an instantiable type to a non-instantiable type only if the type has no columns, views, tables, or instances that reference that type, either directly, or indirectly through another type or subtype.

You cannot declare a non-instantiable type to be `FINAL`, which would be pointless anyway.

## Inheriting, Overloading, and Overriding Methods

A subtype automatically inherits all methods (both member and static methods) declared in or inherited by its supertype.

A subtype can redefine methods it inherits, and it can also add new methods. It can even add new methods that have the same names as methods it inherits, such that the subtype ends up containing more than one method with the same name.

Giving a type multiple methods with the same name is called method overloading. Redefining an inherited method to customize its behavior in a subtype is either overriding, in the case of member methods, or hiding, in the case of static methods.

## Overloading Methods

Overloading is useful when you want to provide a variety of ways of doing something. For example, a shape object might overload a `draw()` method with another `draw()` method that adds a text label to the drawing and contains an argument for the label's text.

When a type has several methods with the same name, the compiler uses the methods' signatures to tell them apart. A method's signature is a sort of structural profile. It consists of the method's name and the number, types, and order of the method's formal parameters (including the implicit `self` parameter). Methods that have the same name but different signatures are called overloads (when they exist in the same type).

In the following example, Subtype `MySubType_typ` creates an overload of `draw()`:

```
CREATE TYPE MyType_typ AS OBJECT (... ,
    MEMBER PROCEDURE draw(x NUMBER), ...) NOT FINAL;

CREATE TYPE MySubType_typ UNDER MyType_typ (... ,
    MEMBER PROCEDURE draw(x VARCHAR2(20)),
    STATIC FUNCTION bar(...)...
...);
```

`MySubType_typ` contains two versions of `draw()`. One is an inherited version with a `NUMBER` parameter and the other has a `VARCHAR2` parameter.

## Redefining Methods

Overriding and hiding redefine an inherited method to make it do something different in the subtype. For example, a subtype `circular_shape` derived from a shape supertype might override a member method `calculate_area()` to customize it specifically for calculating the area of a circle.

Redefining a method is called overriding when the method that is redefined is a member method; redefining is called hiding when the redefined method is a static method. Overriding and hiding are similar in that, in either case, the version of the method redefined in the subtype eclipses an inherited version of the same name and signature such that the new version is executed instead of the inherited one

whenever an instance of the subtype invokes the method. If the subtype itself has subtypes, these inherit the redefined method instead of the original version.

However, with overriding, the system relies on type information contained in the member method's implicit self argument to dynamically choose the correct version of the method to execute. With hiding, the correct version can be identified at compile time, and dynamic dispatch is not necessary. See ["Dynamic Method Dispatch"](#) on page 2-23.

It is possible that a supertype may contain overloads of a method that is redefined in a subtype. Overloads of a method all have the same name, so the compiler uses the signature of the subtype's method to identify the particular version in the supertype that is superseded. This means that, to override or hide a method, you must preserve its signature.

A subtype that overrides a member method must signal the override with the `OVERRIDING` keyword in the type definition. No such special keyword is required when a subtype hides a static method.

For example, in the following code, the subtype signals that it is overriding method `Print()`:

```
CREATE TYPE MyType_typ AS OBJECT (...,  
    MEMBER PROCEDURE Print(),  
    FINAL MEMBER FUNCTION function_mytype(x NUMBER)...  
    ) NOT FINAL;  
  
CREATE TYPE MySubType_typ UNDER MyType_typ (...,  
    OVERRIDING MEMBER PROCEDURE Print(),  
    ...);
```

As with new methods, you supply the declaration for a method that hides or overrides in a `CREATE TYPE BODY` statement.

### Restrictions on Overriding Methods

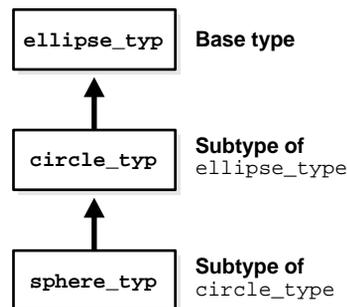
- You can override only methods that are not declared to be `final` in the supertype.
- Order methods may appear only in the root type of a type hierarchy: they may not be redefined (overridden) in subtypes.
- A static method in a subtype may not redefine a member method in the supertype.

- A member method in a subtype may not redefine a static method in the supertype.
- If a method being overridden provides default values for any parameters, then the overriding method must provide the same default values for the same parameters.

## Dynamic Method Dispatch

As a result of method overriding, a type hierarchy can define multiple implementations of the same method. For example, in a hierarchy of the types `ellipse_typ`, `circle_typ`, `sphere_typ`, each type might define a method `calculate_area()` differently.

**Figure 2–2 Hierarchy of Types**



When such a method is invoked, the type of the object instance that invokes it is used to determine which implementation of the method to use. The call is then dispatched to that implementation for execution. This process of selecting a method implementation is called virtual or dynamic method dispatch because it is done at run time, not at compile time.

A method call is dispatched to the nearest implementation, working back up the inheritance hierarchy from the current or specified type. If the call invokes a member method of an object instance, the type of that instance is the current type, and the implementation defined or inherited by that type is used. If the call invokes a static method of a type, the implementation defined or inherited by that specified type is used.

For example, if `c1` is an object instance of `circle_typ`, `c1.proc()` looks first for an implementation of `proc()` defined in `circle_typ`. If none is found, it looks up the supertype chain for an implementation in `ellipse_typ`. The fact that

`sphere_typ` also defines an implementation is irrelevant because the type hierarchy is searched only upwards, toward the top. Subtypes of the current type are not searched.

Similarly, a call to a static method `circle_typ.bar()` looks first in `circle_typ` and then, if necessary, in the supertype(s) of `circle_typ`. The subtype `sphere_typ` is not searched.

**See Also:** *PL/SQL User's Guide and Reference* for information on how subprograms calls are resolved and the dynamic dispatch feature

## Substituting Types in a Type Hierarchy

In a type hierarchy, the subtypes are variant kinds of the root, base type. For example, a `student_typ` type and an `employee_typ` are kinds of a `person_typ`. The base type includes these other types.

When you work with types in a type hierarchy, sometimes you want to work at the most general level and, for example, select or update all persons. But sometimes you want to select or update only students, or only persons who are not students.

The (polymorphic) ability to select all persons and get back not only objects whose declared type is `person_typ` but also objects whose declared (sub)type is `student_typ` or `employee_typ` is called substitutability. A supertype is substitutable if one of its subtypes can substitute or stand in for it in a slot (a variable, column, and so forth) whose declared type is the supertype.

In general, types are substitutable. Object attributes, collection elements and `REFs` are substitutable. An attribute defined as a `REF`, type, or collection of type `person_typ` can hold a `REF` to, an instance of, or instances of an instance of `person_typ` or an instance of any subtype of `person_typ`.

This is what you would expect, given that a subtype is, after all, just a specialized kind of any of its supertypes. Formally, though, a subtype is a type in its own right: it is not the same type as its supertype. A column that holds all persons, including all persons who are students and all persons who are employees, actually holds data of multiple types.

Substitutability comes into play in attributes, columns, and rows (namely, of an object view or object table) declared to be an object type, a `REF` to an object type, or a collection type.

In principle, object attributes, collection elements and `REFs` are always substitutable: there is no syntax at the level of the type definition to constrain their

substitutability to some subtype. You can, however, turn off or constrain substitutability at the storage level, for specific tables and columns.

**See Also:**

- ["Turning Off Substitutability in a New Table"](#) on page 2-29
- ["Constraining Substitutability"](#) on page 2-30

## Column and Row Substitutability

Object type columns are substitutable, and so are object-type rows in object tables and views. In other words, a column or row defined to be of type  $\tau$  can contain instances of  $\tau$  and any of its subtypes.

For example, consider the `person_typ` type hierarchy introduced in ["Creating Subtypes"](#) on page 2-18.

An object table of `person_typ` can contain rows of all three types. You insert an instance of a given type using the constructor for that type in the `VALUES` clause of the `INSERT` statement as shown in [Example 2-1](#).

**Example 2-1 Inserting Values into Substitutable Columns**

```
CREATE TABLE person_obj_table OF person_typ;

INSERT INTO person_obj_table
VALUES (person_typ(12, 'Bob Jones', '111-555-1212'));

INSERT INTO person_obj_table
VALUES (student_typ(51, 'Joe Lane', '1-800-555-1312', 12, 'HISTORY'));

INSERT INTO person_obj_table
VALUES (part_time_student_typ(52, 'Kim Patel', '1-800-555-1232', 14,
'PHYSICS', 20));
```

Similarly, in a relational table or view, a substitutable column of type `person_typ` can contain instances of all three types. The following example inserts a person, a student, and a part-time student in the `person_typ` column `contact`:

```
CREATE TABLE contacts (
    contact          person_typ,
    contact_date     DATE );

INSERT INTO contacts
VALUES (person_typ (12, 'Bob Jones', '111-555-1212'), '24 Jun 2003' );
```

```
INSERT INTO contacts
VALUES (student_typ(51, 'Joe Lane', '1-800-555-1312', 12, 'HISTORY'),
       '24 Jun 2003' );
```

```
INSERT INTO contacts
VALUES (part_time_student_typ(52, 'Kim Patel', '1-800-555-1232', 14,
                              'PHYSICS', 20), '24 Jun 2003' );
```

A newly created subtype can be stored in any substitutable tables and columns of its supertype, including tables and columns that existed before the subtype was created.

Attributes in general can be accessed using the dot notation. Attributes of a subtype of a row or column's declared type can be accessed with the `TREAT` function. For example:

```
SELECT TREAT(contact AS student_typ).major FROM contacts;
```

**See Also:** ["TREAT"](#) on page 2-39

### Using OBJECT\_VALUE and OBJECT\_ID with Substitutable Rows

The `OBJECT_VALUE` and `OBJECT_ID` pseudocolumns allow you to access and identify the value and OID of a substitutable row in an object table.

In [Example 2-2](#), the `categories_tab` is a table of `category_typ` objects. The `composite_category_typ` and `leaf_category_typ` types are subtypes of `category_typ`. Both subtypes contain the `category_name`, `category_description`, and `category_id` attributes of the `category_typ` supertype. Additionally, `leaf_category_typ` contains the `product_ref_list_typ` nested table and `composite_category_typ` contains the `subcategory_ref_list` nested table.

#### **Example 2-2 Using OBJECT\_VALUE and OBJECT\_ID**

```
CREATE TABLE categories_tab OF category_typ
( category_id PRIMARY KEY)
  NESTED TABLE TREAT
  (OBJECT_VALUE AS leaf_category_typ).product_ref_list
  STORE AS product_ref_list_nestedtab
  NESTED TABLE TREAT
  (OBJECT_VALUE AS composite_category_typ).subcategory_ref_list
  STORE AS subcategory_ref_list_nestedtab;
```

```
SELECT c.category_id, o.*
   FROM categories_tab c,
        TABLE (TREAT(c.OBJECT_VALUE AS leaf_category_typ).product_ref_list) o
 WHERE c.category_id = 31;
```

```
SELECT c.OBJECT_ID, c.category_id, o.COLUMN_VALUE.category_id
   FROM categories_tab c,
        TABLE (TREAT(c.OBJECT_VALUE AS composite_category_typ).subcategory_ref_list) o
 WHERE category_id = 10;
```

## Subtypes Having Supertype Attributes

A subtype can have an attribute that is a supertype. For example:

```
CREATE TYPE student_typ UNDER person_typ (
    dept_id    NUMBER,
    major      VARCHAR2(30),
    advisor    person_typ);
/
```

However, columns of such types are not substitutable. Similarly, a subtype `ST` can have a collection attribute whose element type is one of `ST`'s supertypes, but, again, columns of such types are not substitutable. For example, if `student_typ` had a nested table or varray of `person_typ`, the `student_typ` column would not be substitutable.

You can, however, define substitutable columns of subtypes that have `REF` attributes that reference supertypes. For example, the `composite_category_typ` subtype shown in [Example 2-2](#) contains the `subcategory_ref_list` nested table. This table contains `subcategory_ref_list_typ` which are `REFs` to `category_typ`. The subtype was created as follows:

```
CREATE TYPE subcategory_ref_list_typ
  AS TABLE OF REF category_typ;
/

CREATE TYPE composite_category_typ
  UNDER category_typ
  (
    subcategory_ref_list subcategory_ref_list_typ
  ...
```

**See Also:** ["Turning Off Substitutability in a New Table"](#) on page 2-29

## REF Columns and Attributes

REF columns and attributes are substitutable in both views and tables. For example, in either a view or a table, a column declared to be REF `person_typ` can hold references to instances of `person_typ` or any of its subtypes.

## Collection Elements

Collection elements are substitutable in both views and tables. For example, a nested table of `person_typ` can contain object instances of `person_typ` or any of its subtypes.

## Creating Subtypes After Creating Substitutable Columns

If you create a subtype, any table that already has substitutable columns of the supertype is automatically enabled to store the new subtype as well. This means that your options for creating subtypes are affected by the existence of such tables. If such a table exists, you can only create subtypes that are substitutable, that is, subtypes that Oracle can enable that table to store.

The following example shows an attempt to create a subtype `student_typ` under `person_typ`.

```
CREATE TYPE person_typ AS OBJECT (  
    idno          NUMBER,  
    name          VARCHAR2(30),  
    phone         VARCHAR2(20))  
    NOT FINAL;  
/  
  
CREATE TYPE employee_typ UNDER person_typ (  
    emp_id NUMBER,  
    mgr VARCHAR2(30))  
    NOT FINAL;  
/  
  
CREATE TABLE person_obj_table (p person_typ);
```

The following statement fails because `student_typ` has a supertype attribute, and table `person_obj_table` has a substitutable column `p` of the supertype.

```
CREATE TYPE student_typ UNDER person_typ (  
    advisor person_typ);  
/
```

The following attempt succeeds. This version of the `student_typ` subtype is substitutable. Oracle automatically enables table `person_obj_table` to store instances of this new type.

```
CREATE TYPE student_typ UNDER person_typ (
    dept_id NUMBER,
    major VARCHAR2(30));
/
INSERT INTO person_obj_table
VALUES (student_typ(51, 'Joe Lane', '1-800-555-1312', 12, 'HISTORY'));
```

## Dropping Subtypes After Creating Substitutable Columns

You can drop a subtype with the `VALIDATE` option only if no instances of the subtype are stored in any substitutable column of the supertype.

For example, the following statement fails because an instance of `student_typ` is stored in substitutable column `p` of table `person_obj_table`:

```
DROP TYPE student_typ VALIDATE;
```

To drop the type, first delete any of its instances in substitutable columns of the supertype:

```
DELETE FROM person_obj_table WHERE p IS OF (student_typ);

DROP TYPE student_typ VALIDATE;
```

## Turning Off Substitutability in a New Table

When creating a table, you can turn off all substitutability on a column or attribute, including embedded attributes and collections nested to any level, with the clause `NOT SUBSTITUTABLE AT ALL LEVELS`.

In the following example, the clause confines column `office` of a relational table to storing only `office_typ` instances and disallows any subtype instances:

```
CREATE TYPE office_typ AS OBJECT (
    office_id    VARCHAR(10),
    location     location_typ,
    occupant     person_typ )
NOT FINAL;
/

CREATE TABLE dept_office (
    dept_no     NUMBER,
```

```
office      office_typ)
COLUMN office NOT SUBSTITUTABLE AT ALL LEVELS;
```

With object tables, the clause can be applied to the table as a whole, like this:

```
CREATE TABLE office_tab OF office_typ
  NOT SUBSTITUTABLE AT ALL LEVELS;
```

Alternatively, the clause can also be applied to turn off substitutability in a particular column that is, for a particular attribute of the object type of the table:

```
CREATE TABLE office_tab OF office_typ
  COLUMN occupant NOT SUBSTITUTABLE AT ALL LEVELS;
```

You can specify that the element type of a collection is not substitutable using syntax like the following:

```
CREATE TABLE people_tab (
  people_column people_typ )
  NESTED TABLE people_column
    NOT SUBSTITUTABLE AT ALL LEVELS STORE AS people_column_nt;
```

There is no mechanism to turn off substitutability for REF columns.

## Constraining Substitutability

You can impose a constraint that limits the range of subtypes permitted in an object column or attribute to a particular subtype in the declared type's hierarchy. You do this using an *IS OF type* constraint.

For example, the following statement creates a table of *office\_typ* in which occupants are constrained to just those persons who are employees:

```
CREATE TABLE office_tab OF office_typ
  COLUMN occupant IS OF (ONLY employee_typ);
```

Although the type *office\_typ* allows authors to be of type *person\_typ*, the column declaration imposes a constraint to store only instances of *employee\_typ*.

You can only use the *IS OF type* operator to constrain row and column objects to a single subtype (not several), and you must use the *ONLY* keyword, as in the preceding example.

You can use either *IS OF type* or *NOT SUBSTITUTABLE AT ALL LEVELS* to constrain an object column, but you cannot use both.

## Modifying Substitutability

In an existing table, you can change an object column from `SUBSTITUTABLE` to `NOT SUBSTITUTABLE` (or from `NOT SUBSTITUTABLE` to `SUBSTITUTABLE`) by using an `ALTER TABLE` statement. To do so, you specify the clause `[NOT] SUBSTITUTABLE AT ALL LEVELS` for the particular column.

You can modify substitutability only for a specific column; you cannot modify substitutability for an object table as a whole.

The following statement makes column `office` substitutable:

```
ALTER TABLE dept_office
  MODIFY COLUMN office SUBSTITUTABLE AT ALL LEVELS;
```

The following statement makes the column not substitutable. Notice that it also uses the `FORCE` keyword. This keyword causes any hidden columns containing typeid information or data for subtype attributes to be dropped:

```
ALTER TABLE dept_office
  MODIFY COLUMN office NOT SUBSTITUTABLE AT ALL LEVELS FORCE;
```

If the `FORCE` keyword is not used when a column is made not substitutable, the column and all attributes of the type must be `FINAL` or the `ALTER TABLE` statement will fail.

A `VARRAY` column can be modified from `SUBSTITUTABLE` to `NOT SUBSTITUTABLE` only if the element type of the varray is `final` itself and has no embedded types (in its attributes or in their attributes, and so on) that are not `final`.

See "[Hidden Columns for Substitutable Columns and Tables](#)" on page 7-3 for more information about hidden columns for typeids and subtype attributes.

## Restrictions on Modifying Substitutability

You can change the substitutability of only one column at a time with an `ALTER TABLE` statement. To change substitutability for multiple columns, you must issue multiple statements.

In an object table, you can modify substitutability for a column only if substitutability was not explicitly set at the table level, for the entire table, when the table was created.

For example, the following attempt to modify substitutability for column `address` succeeds because substitutability has not been explicitly turned on or off at the table level in the `CREATE TABLE` statement:

```
CREATE TABLE office_tab OF office_typ;

ALTER TABLE office_tab
  MODIFY COLUMN occupant NOT SUBSTITUTABLE AT ALL LEVELS FORCE;
```

However, in the following example, substitutability is explicitly set at the table level, so the attempt to modify the setting for column address fails:

```
CREATE TABLE office_tab OF office_typ
  NOT SUBSTITUTABLE AT ALL LEVELS;

/* Following SQL statement generates an error: */
ALTER TABLE office_tab
  MODIFY COLUMN occupant SUBSTITUTABLE AT ALL LEVELS FORCE;
```

A column whose substitutability is already constrained by an `IS OF type` operator cannot have its substitutability modified with a `[NOT] SUBSTITUTABLE AT ALL LEVELS` clause. See "[Constraining Substitutability](#)" on page 2-30 for information about `IS OF type`.

## Assignments Across Types

The assignment rules described in this section apply to `INSERT/UPDATE` statements, the `RETURNING` clause, function parameters, and `PL/SQL` variables.

### Objects and REFs to Objects

Substitutability is the ability of a subtype to stand in for one of its supertypes. An attempt to perform a substitution in the other direction, to substitute a supertype for a subtype, raises an error at compile time.

An assignment of a source of type `source_typ` to a target of type `target_typ` must be of one of the following two patterns:

- Case 1: `source_typ` and `target_typ` are the same type
- Case 2: `source_typ` is a subtype of `target_typ` (widening)

Case 2 illustrates widening. Widening is an assignment in which the declared type of the source is more specific than the declared type of the target. For example, assigning an employee instance to a variable of person type.

Intuitively, the idea here is that you are regarding an employee as a person. An employee is a more narrowly defined, specialized kind of person, so you can put an employee in a slot meant for a person if you do not mind ignoring whatever extra

specialization makes that person an employee. All employees are persons, so a widening assignment always works.

To illustrate widening, suppose that you have the following table:

```
TABLE T(pers_col person_typ, emp_col employee_typ, stu_col student_typ)
```

The following assignments show widening. The assignments are valid unless perscol has been defined to be not substitutable.

```
UPDATE T set pers_col = emp_col;
```

The following is a PL/SQL example:

```
declare
  var1 person_typ;
  var2 employee_typ;
begin
  var1 := var2;
end;
```

Besides widening, there is also narrowing. Narrowing is the reverse of widening. It involves regarding a more general, less specialized type of thing, such as a person, as a more narrowly defined type of thing, such as an employee. Not all persons are employees, so a particular assignment like this works only if the person in question actually happens to be an employee.

To do a narrowing assignment, you must use the `TREAT` function to test that the source instance of the more general declared type is in fact an instance of the more specialized target type and can therefore be operated on as such. The `TREAT` function does a runtime check to confirm this and returns `NULL` if the source value the person in question is not of the target type or one of its subtypes.

For example, the following `UPDATE` statement sets values of `person_typ` in column `perscol` into column `empcol` of `employee_typ`. For each value in `perscol`, the assignment succeeds only if that person is also an employee. If person George is not an employee, `TREAT` returns `NULL`, and the assignment returns `NULL`.

```
UPDATE T set emp_col = TREAT(pers_col AS employee_typ);
```

The following statement attempts to do a narrowing assignment without explicitly changing the declared type of the source value. The statement will return an error:

```
UPDATE T set emp_col = pers_col;
```

**See Also:** ["TREAT"](#) on page 2-39

## Collection Assignments

In assignments of expressions of a collection type, the source and target must be of the same declared type. Neither widening nor narrowing is permitted. However, a subtype value can be assigned to a supertype collection. For example, suppose we have the following collection types:

```
CREATE TYPE person_set AS TABLE OF person_typ;
/

CREATE TYPE student_set AS TABLE OF student_typ;
/
```

Expressions of these different collection types cannot be assigned to each other, but a collection element of `student_typ` can be assigned to a collection of `PersonSet` type:

```
DECLARE
    var1 person_set;
    var2 student_set;
    elem1 person_typ;
    elem2 student_typ;
begin
    var1 := var2;           /* ILLEGAL - collections not of same type */
    var1 := person_set (elem1, elem2); /* LEGAL : Element is of subtype */
```

## Comparisons of Objects, REF Variables, and Collections

This section discusses the comparison operators used in SQL conditions.

**See Also:** *Oracle Database SQL Reference* for information about using SQL conditions

### Comparing Object Instances

Two object instances can be compared if, and only if, they are both of the same declared type, or one is a subtype of the other.

Map methods and order methods provide the mechanism for comparing objects. You optionally define one or the other of these in an object type to specify the basis on which you want objects of that type to be compared. If a method of either sort is defined, it is called automatically whenever objects of that type or one of its subtypes need to be compared.

If a type does not define either a map method or an order method, object variables of that type can be compared only in SQL statements and only for equality or inequality. Two objects of the same type count as equal only if the values of their corresponding attributes are equal.

**See Also:** ["Methods for Comparing Objects"](#) on page 2-10

### Comparing REF Variables

Two REF variables can be compared if, and only if, the targets that they reference are both of the same declared type, or one is a subtype of the other.

## Functions and Operators Useful with Objects

Several functions and operators are particularly useful for working with objects and references to objects:

- [CAST](#)
- [CURSOR](#)
- [DEREF](#)
- [IS OF type](#)
- [REF](#)
- [SYS\\_TYPEID](#)
- [TABLE\(\)](#)
- [TREAT](#)
- [VALUE](#)

Examples are given throughout this book.

In PL/SQL the VALUE, REF and Deref functions can appear only in a SQL statement.

**See Also:** *Oracle Database SQL Reference* for information about SQL functions

### CAST

CAST converts one built-in datatype or collection-typed value into another built-in datatype or collection-typed value.

For example:

```
CREATE TYPE person_list_typ AS TABLE OF person_typ;
/

SELECT CAST(COLLECT(contact) AS person_list_typ)
FROM contacts;
```

**See Also:** *Oracle Database SQL Reference*

## CURSOR

A **CURSOR** expression returns a nested cursor. This form of expression is equivalent to the **PL/SQL REF CURSOR** and can be passed as a **REF CURSOR** argument to a function.

**See Also:** *Oracle Database SQL Reference*

## DEREF

The **DEREF** function in a SQL statement returns the object instance corresponding to a **REF**. The object instance returned by **DEREF** may be of the declared type of the **REF** or any of its subtypes.

For example, the following statement returns **person\_typ** objects from the table **contact\_ref**.

```
SELECT Deref(c.contact_ref), c.contact_date
FROM contacts_ref c;
```

See "[Dereferencing REFs](#)" on page 1-11.

## IS OF *type*

The **IS OF *type*** predicate tests object instances for the level of specialization of their type.

For example, the following query retrieves all student instances (including any subtypes of students) stored in the **person\_obj\_table** table.

```
SELECT VALUE(p)
FROM person_obj_table p
WHERE VALUE(p) IS OF (student_typ);
```

For any object that is not of a specified subtype, or a subtype of a specified subtype, `IS OF` returns `FALSE`. Subtypes of a specified subtype are just more specialized versions of the specified subtype. If you want to exclude such subtypes, you can use the `ONLY` keyword. This keyword causes `IS OF` to return `FALSE` for all types except the specified types.

In the following example, the statement tests objects in object table `person_obj_table`, which contains persons, employees, and students, and returns `REFs` just to objects of the two specified person subtypes `employee_typ`, `student_typ`, and their subtypes, if any:

```
SELECT REF(p)
       FROM person_obj_table p
WHERE VALUE(p) IS OF (employee_typ, student_typ);
```

Here is a similar example in PL/SQL. The code does something if the person is an employee or student:

```
DECLARE
    var person_t;
BEGIN
    if var is of (employee_t, student_t) then
        ...
    END;
```

The following statement returns only students whose most specific or specialized type is `student_typ`. If the table or view contains any objects of a subtype of `student_typ`, such as `part_time_student_typ`, these are excluded. The example uses the `TREAT` function to convert objects that are students to `student_typ` from the declared type of the view, `person_typ`:

```
SELECT TREAT(VALUE(p) AS student_typ)
       FROM person_obj_table p
WHERE VALUE(p) IS OF(ONLY student_typ);
```

To test the type of the object that a `REF` points to, you can use the `DEREF` function to dereference the `REF` before testing with the `IS OF type` predicate.

For example, if `contact_ref` is declared to be `REF person_typ`, you can get just the rows for students as follows:

```
SELECT *
       FROM contacts_ref
WHERE DEREF(contact_ref) IS OF (student_typ);
```

## REF

The `REF` function in a SQL statement takes as an argument a correlation name for an object table or view and returns a reference (a `REF`) to an object instance from that table or view. The `REF` function may return references to objects of the declared type of the table, view, or any of its subtypes. For example, the following statement returns the references to all persons, including references to students and employees, whose `idno` attribute is 12:

```
SELECT REF(p)
   FROM person_obj_table p
  WHERE p.idno = 12;
```

## SYS\_TYPEID

The `SYS_TYPEID` function can be used in a query to return the typeid of the most specific type of the object instance passed to the function as an argument.

The most specific type of an object instance is the type to which the instance belongs that is farthest removed from the root type. For example, if Tim is a part-time student, he is also a student and a person, but his most specific type is part-time student.

The function returns the typeids from the hidden type discriminant column that is associated with every substitutable column. The function returns a null typeid for a final, root type.

The syntax of the function is:

```
SYS_TYPEID( object_type_value )
```

Function `SYS_TYPEID` may be used only with arguments of an object type. Its primary purpose is to make it possible to build an index on a hidden type discriminant column.

All types that do belong to a type hierarchy are assigned a non-null typeid that is unique within the type hierarchy. Types that do not belong to a type hierarchy have a null typeid.

Every type except a final root type belongs to a type hierarchy. A final root type has no types related to it by inheritance:

- It cannot have subtypes derived from it because it is final
- It is not itself derived from some other type because it is a root type, so it does not have any supertypes.

**See Also:** ["Hidden Columns for Substitutable Columns and Tables"](#) on page 7-3 for more information about type discriminant columns

For an example of `SYS_TYPEID`, consider the substitutable object table `person_obj_table`, of `person_typ`. `person_typ` is the root type of a hierarchy that has `student_typ` as a subtype and `part_time_student_typ` as a subtype of `student_typ`. See [Example 2-1](#) on page 2-25.

The following query uses `SYS_TYPEID`. It gets the name attribute and typeid of the object instances in the `person_obj_table` table. Each of the instances is of a different type:

```
SELECT name, SYS_TYPEID(VALUE(p)) typeid FROM person_obj_table p;
```

**See Also:** ["Hidden Columns for Substitutable Columns and Tables"](#) on page 7-3 for information about the type discriminant and other hidden columns

## TABLE()

Table functions are functions that produce a collection of rows, a nested table or a varray, that can be queried like a physical database table or assigned to a PL/SQL collection variable. You can use a table function like the name of a database table, in the `FROM` clause of a query, or like a column name in the `SELECT` list of a query.

A table function can take a collection of rows as input. An input collection parameter can be either a collection type, such as a `VARRAY` or a PL/SQL table, or a `REF CURSOR`.

Use `PIPELINED` to instruct Oracle to return the results of a table function iteratively. A table function returns a nested table or varray collection type. You query table functions by using the `TABLE` keyword before the function name in the `FROM` clause of the query.

**See Also:** *PL/SQL User's Guide and Reference* for information on `TABLE()` functions

## TREAT

The `TREAT` function does a runtime check to confirm that an expression can be operated on as if it were of a different specified type in the hierarchy normally, a subtype of the expression's declared type. In other words, the function attempts to

treat a supertype instance as a subtype instance to treat a person as a student, for example. Whether this can be done in a given case depends on whether the person in question actually is a student (or student subtype, such as a part-time student). If the person is a student, then the person is returned as a student, with the additional attributes and methods that a student may have. If the person happens not to be a student, `TREAT` returns `NULL` in SQL.

The two main uses of `TREAT` are:

- In narrowing assignments, to modify the type of an expression so that the expression can be assigned to a variable of a more specialized type in the hierarchy: in other words, to set a supertype value into a subtype.
- To access attributes or methods of a subtype of the declared type of a row or column

A substitutable object table or column of type `T` has a hidden column for every attribute of every subtype of `T`. These hidden columns are not listed by a `DESCRIBE` statement, but they contain subtype attribute data. `TREAT` enables you to access these columns.

The following example shows `TREAT` used in an assignment where a column of person type is set into a column of employee type. For each row in `perscol`, `TREAT` returns an employee type or `NULL`, depending on whether the given person happens to be an employee.

```
UPDATE T set empcol = TREAT(perscol AS employee_typ);
```

In the next example, `TREAT` returns all (and only) `student_typ` instances from `person_obj_table` of type `person_typ`, a supertype of `student_typ`. The statement uses `TREAT` to modify the type of `p` from `person_typ` to `student_typ`.

```
SELECT TREAT(VALUE(p) AS student_typ)
FROM person_obj_table p;
```

For each `p`, the `TREAT` modification succeeds only if the most specific or specialized type of the value of `p` is `student_typ` or one of its subtypes. If `p` is a person who is not a student, or if `p` is `NULL`, `TREAT` returns `NULL` in SQL or, in PL/SQL, raises an exception.

You can also use `TREAT` to modify the declared type of a `REF` expression. For example:

```
SELECT TREAT(REF(p) AS REF student_typ)
FROM person_obj_table p;
```

The previous example returns `REFs` to all `student_typ` instances. In SQL it returns `NULL REFs` for all person instances that are not students, and in PL/SQL it raises an exception.

Perhaps the most important use of `TREAT` is to access attributes or methods of a subtype of a row or column's declared type. The following query retrieves the `major` attribute of all persons, students and part-time students, who have this attribute. `NULL` is returned for persons who are not students:

```
SELECT name, TREAT(VALUE(p) AS student_typ).major major
       FROM person_obj_table p;
```

The following query will not work because `major` is an attribute of `student_typ` but not of `person_typ`, the declared type of table `persons`:

```
SELECT name, VALUE(p).major major
       FROM person_obj_table p;
```

The following is a PL/SQL example:

```
DECLARE
    var Person_t;
BEGIN
    TREAT(var as Employee_t).raise_salary;
END;
```

**See Also:** ["Assignments Across Types"](#) on page 2-32 for information on using `TREAT` in assignments.

## VALUE

In a SQL statement, the `VALUE` function takes as its argument a correlation variable (table alias) for an object table or object view and returns object instances corresponding to rows of the table or view. The `VALUE` function may return instances of the declared type of the row or any of its subtypes. For example, the following query returns all persons, including students and employees, from table `person_obj_table` of `person_typ`:

```
SELECT VALUE(p) FROM person__obj_table p;
```

To retrieve only part time students, that is, instances whose most specific type is `part_time_student_typ`, use the `ONLY` keyword to confine the selection:

```
SELECT VALUE(p) FROM person_obj_table p
       WHERE VALUE(p) IS OF (ONLY part_time_student_typ);
```

The following example shows `VALUE` used to return object instance rows for updating:

```
UPDATE TABLE(SELECT d.dept_emps FROM department_persons d
               WHERE d.dept_no = 101) p
SET VALUE(p) = person_typ(2, 'Diane Smith', '1-800-555-1243')
WHERE p.idno = 2;
```

---

## Support for Collection Datatypes

This chapter provides basic information about working with objects. It explains what object types, methods, and collections are and describes how to create and work with a hierarchy of object types that are derived from a shared root type and are connected by inheritance.

This chapter contains these topics:

- [Creating Collection Datatypes](#)
- [Operations on Collection Datatypes](#)

## Creating Collection Datatypes

Oracle supports the varray and nested table collection datatypes.

- A varray is an ordered collection of elements
- A nested table can have any number of elements

If you need to store only a fixed number of items, or to loop through the elements in order, or you will often want to retrieve and manipulate the entire collection as a value, then use a varray.

If you need to run efficient queries on a collection, handle arbitrary numbers of elements, or do mass insert/update/delete operations, then use a nested table.

**See Also:** ["Design Considerations for Collections"](#) on page 8-11.

## Creating an Instance of a VARRAY or Nested Table

You create an instance of a collection type in the same way that you create an instance of any other object type, namely, by calling the type's constructor method. The name of a type's constructor method is simply the name of the type. You specify the elements of the collection as a comma-delimited list of arguments to the method.

Calling a constructor method with an empty list creates an empty collection of that type. Note that an empty collection is an actual collection that happens to be empty; it is not the same as a null collection.

**See Also:**

- *Oracle Database SQL Reference* for information about nested cursors
- ["Design Considerations for Nested Tables"](#) on page 8-14 for more information on using nested tables

## Constructor Methods for Collections

The following example illustrates a literal invocation of the constructor method for the nested table type `people_ttyp`.

```
CREATE TYPE people_ttyp AS TABLE OF person_ttyp;
/
people_ttyp ( person_ttyp(1, 'John Smith', '1-800-555-1212'),
              person_ttyp(2, 'Diane Smith', NULL) )
```

The following examples show how it is used in SQL statements to insert values into a nested table.

```
CREATE TABLE people_tab (
    group_no NUMBER,
    people_column people_typ )
    NESTED TABLE people_column STORE AS people_column_nt;

INSERT INTO people_tab VALUES (
    100,
    people_typ( person_typ(1, 'John Smith', '1-800-555-1212'),
               person_typ(2, 'Diane Smith', NULL)));
```

When you declare a table column to be of an object type or collection type, you can include a `DEFAULT` clause. This provides a value to use in cases where you do not explicitly specify a value for the column. The default clause must contain a literal invocation of the constructor method for that object or collection.

The following example shows how to use literal invocations of constructor methods to specify defaults:

**Example 3-1 Creating the `department_persons` Table**

```
CREATE TABLE department_persons (
    dept_no    NUMBER PRIMARY KEY,
    dept_name  CHAR(20),
    dept_mgr   person_typ DEFAULT person_typ(10, 'John Doe', NULL),
    dept_emps  people_typ DEFAULT people_typ() )
    NESTED TABLE dept_emps STORE AS dept_emps_tab;
```

Note that the term `people_typ()` is a literal invocation of the constructor method for an empty `people_typ` table.

## Varrays

A varray is an ordered set of data elements. All elements of a given varray are of the same datatype or a subtype of the declared one. Each element has an index, which is a number corresponding to the element's position in the array. The index number is used to access a specific element.

When you define a varray, you specify the maximum number of elements it can contain, although you can change this number later. The number of elements in an array is the size of the array. Oracle allows arrays to be of variable size, which is why they are called varrays.

For example, the following statement declares an array type:

```
CREATE TYPE email_list_arr AS VARRAY(10) OF VARCHAR2(80);  
/
```

The `VARRAY`s of type `email_list_arr` have no more than ten elements, each of datatype `VARCHAR2(80)`.

Creating an array type, as with a SQL object type, does not allocate space. It defines a datatype, which you can use as:

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

A varray is normally stored in line, that is, in the same tablespace as the other data in its row. If it is sufficiently large, Oracle stores it as a `BLOB`. See "[Storage Considerations for Varrays](#)" on page 8-13.

You can create a `VARRAY` type of `XMLType` or of a `LOB` type for procedural purposes, such as in PL/SQL or in view queries. However, database storage of a varray of those types is not supported. This means that you cannot create an object table or an object type column of a varray type of `XMLType` or of a `LOB` type.

**See Also:** *Oracle Database SQL Reference* for information and examples on the `STORE AS LOB` clause of the `CREATE TABLE` statement

## Nested Tables

A nested table is an unordered set of data elements, all of the same datatype. No maximum is specified in the definition of the table and the order of the elements is not preserved. You select, insert, delete, and update in a nested table just as you do with ordinary tables using the `TABLE` expression.

Elements of a nested table are actually stored in a separate storage table that contains a column that identifies the parent table row or object to which each element belongs. A nested table has a single column, and the type of that column is a built-in type or an object type. If the column in a nested table is an object type, the table can also be viewed as a multi-column table, with a column for each attribute of the object type.

In [Example 3-2](#) on page 3-5, the table type used for the nested tables is declared with the `CREATE TYPE ... IS TABLE OF` statement. A table type definition 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

When a column in a relational table is of nested table type, Oracle stores the nested table data for all rows of the relational table in the same storage table. Similarly, with an object table of a type that has a nested table attribute, Oracle stores nested table data for all object instances in a single storage table associated with the object table. In [Example 3-2](#), the `NESTED TABLE` clause specifies the storage name for the nested table. The example uses `person_typ` defined in [Example 1-1](#) on page 1-4.

**See Also:** See [Figure 8-2, "Nested Table Storage"](#) on page 8-15

### **Example 3-2** *Creating and Populating Simple Nested Tables*

```
CREATE TYPE people_typ AS TABLE OF person_typ;
/
CREATE TABLE students (
  graduation DATE,
  math_majors people_typ,
  chem_majors people_typ,
  physics_majors people_typ)
NESTED TABLE math_majors STORE AS math_majors_nt
NESTED TABLE chem_majors STORE AS chem_majors_nt
NESTED TABLE physics_majors STORE AS physics_majors_nt;

INSERT INTO students (graduation) VALUES ('01-JUN-03');
UPDATE students
  SET math_majors =
      people_typ (person_typ(12, 'Bob Jones', '111-555-1212'),
                 person_typ(31, 'Sarah Chen', '111-555-2212'),
                 person_typ(45, 'Chris Woods', '111-555-1213')),
      chem_majors =
      people_typ (person_typ(51, 'Joe Lane', '111-555-1312'),
                 person_typ(31, 'Sarah Chen', '111-555-2212'),
                 person_typ(52, 'Kim Patel', '111-555-1232')),
      physics_majors =
      people_typ (person_typ(12, 'Bob Jones', '111-555-1212'),
                 person_typ(45, 'Chris Woods', '111-555-1213'))
WHERE graduation = '01-JUN-03';
```

A convenient way to access the elements of a nested table individually is to use a nested cursor.

## Specifying a Tablespace When Storing a Nested Table

A nested table can be stored in a different tablespace than its parent table. In the following SQL statement, the nested table is stored in the `users` tablespace:

```
CREATE TABLE people_tab (  
    people_column people_typ )  
    NESTED TABLE people_column STORE AS people_column_nt (TABLESPACE users);
```

If the `TABLESPACE` clause is not specified, then the storage table of the nested table is created in the tablespace where the parent table is created. For multilevel nested tables, Oracle creates the child table in the same tablespace as its immediate preceding parent table.

The user can issue `ALTER TABLE MOVE` statement to move a table to a different tablespace. If the user issues `ALTER TABLE MOVE` statement on a table with nested table columns, it only moves parent table, no action is taken on the nested table's storage tables. If the user wants to move a nested table's storage table to a different tablespace, issue `ALTER TABLE MOVE` on the storage table. For example:

```
ALTER TABLE people_tab MOVE TABLESPACE users;  
  
ALTER TABLE people_column_nt MOVE TABLESPACE example;
```

Now the `people_tab` table is in `users` tablespace and nested table is stored in the `example` tablespace.

## Varray Storage

Multilevel varrays are stored in one of two ways, depending on whether the varray is a varray of varrays or a varray of nested tables.

- In a varray of varrays, the entire varray is stored inline in the row unless it is larger than approximately 4000 bytes or LOB storage is explicitly specified.
- In a varray of nested tables, the entire varray is stored in a LOB, with only the LOB locator stored in the row. There is no storage table associated with nested table elements of a varray. The entire nested table collection is stored inside the varray.

You can explicitly specify LOB storage for varrays. The following example does this for the varray elements of a nested table. As the example also shows, you can use the `COLUMN_VALUE` keyword with varrays as well as nested tables.

```
CREATE TYPE email_list_arr AS VARRAY(10) OF VARCHAR2(80);
/

CREATE TYPE email_list_typ AS TABLE OF email_list_arr;
/

CREATE TABLE dept_email_list (
    dept_no NUMBER,
    email_addrs email_list_typ)
    NESTED TABLE email_addrs STORE AS email_addrs_nt
    (VARRAY COLUMN_VALUE STORE AS LOB dept_emails_lob);
```

The following example shows explicit LOB storage specified for a varray of varray type:

```
CREATE TYPE email_list_typ2 AS OBJECT (
    section_no NUMBER,
    emails email_list_arr);
/

CREATE TYPE email_varray_typ AS VARRAY(5) OF email_list_typ2;
/

CREATE TABLE dept_email_list2 (
    dept_no NUMBER,
    email_addrs email_varray_typ)
    VARRAY email_addrs STORE AS LOB dept_emails_lob2;
```

**See Also:** ["Storage Considerations for Varrays"](#) on page 8-13.

## Increasing the Size and Precision of VARRAYs and Nested Tables

When the element type of a `VARRAY` type is a variable character or `RAW` type or a numeric type, you can increase the size of the variable character or `RAW` type or increase the precision of the numeric type. A new type version is generated for the `VARRAY` type.

Options like `INVALIDATE` and `CASCADE` are provided to either invalidate all dependent objects or propagate the change to its type and table dependents

For example:

```
CREATE TYPE email_list_arr AS VARRAY(10) OF VARCHAR2(80);  
/  
  
ALTER TYPE email_list_arr MODIFY ELEMENT TYPE VARCHAR2(100) CASCADE;
```

The same change can be applied to nested table types.

```
CREATE TYPE email_list_tab AS TABLE OF VARCHAR2(30);  
/  
  
ALTER TYPE email_list_tab MODIFY ELEMENT TYPE VARCHAR2(40) CASCADE;
```

## Increasing VARRAY Limit Size

The `ALTER TYPE ... MODIFY LIMIT` syntax allows increasing the number of elements of a `VARRAY` type. If the number of elements of the `VARRAY` type is increased, a new type version is generated for the `VARRAY` type and this is maintained as part of the history of the type changes.

Options like `INVALIDATE` and `CASCADE` are provided to either invalidate all dependent objects or propagate the change to its type and table dependents.

For example:

```
CREATE TYPE email_list_typ AS OBJECT (  
    section_no    NUMBER,  
    emails        email_list_arr);  
/  
  
CREATE TYPE email_varray_typ AS VARRAY(5) OF email_list_typ;  
/  
  
ALTER TYPE email_varray_typ MODIFY LIMIT 100 INVALIDATE;
```

When a `VARRAY` type is altered, changes are propagated to the dependent tables. See ["Propagating VARRAY Size Change"](#) on page 8-13.

## Creating a Varray Containing LOB References

In the following examples, `email_addr`s of type `email_list_typ` already exists in table `dept_email_list` as shown in the SQL examples in ["Varray Storage"](#) on page 3-6.

To create a varray of LOB references, first define a `VARRAY` type of type `REF email_list_typ2`. For example:

```
CREATE TYPE ref_email_varray_typ AS VARRAY(5) OF REF email_list_typ2;
/
```

Next define a column of the array type in `dept_email_list3`.

```
CREATE TABLE dept_email_list3 (
    dept_no NUMBER,
    email_addrs ref_email_varray_typ)
    VARRAY email_addrs STORE AS LOB dept_emails_lob3;
```

## Multilevel Collection Types

Multilevel collection types are collection types whose elements are themselves directly or indirectly another collection type. Possible multilevel collection types are:

- Nested table of nested table type
- Nested table of varray type
- Varray of nested table type
- Varray of varray type
- Nested table or varray of a user-defined type that has an attribute that is a nested table or varray type

Like ordinary, single-level collection types, multilevel collection types can be used as columns in a relational table or with object attributes in an object table.

### Nested Table Storage Tables for Multilevel Collection Types

A nested table type column or object table attribute requires a storage table where rows for all nested tables in the column are stored. With a multilevel nested table collection of nested tables, the inner set of nested tables also requires a storage table just as the outer set does. You specify one by appending a second nested-table storage clause.

[Example 3-3](#) creates a multilevel collection type that is a nested table of nested tables. The example models a system of stars in which each star has a nested table collection of the planets revolving around it, and each planet has a nested table collection of its satellites. In the example, the SQL statements create a table `stars` that contains a column `planets` whose type is a multilevel collection. This multilevel collection is a nested table of an object type that has a nested table attribute `satellites`. Separate nested table clauses are provided for the outer `planets` nested table and for the inner `satellites` one.

**Example 3-3 Nested Table Storage**

```

CREATE TYPE satellite_t AS OBJECT (
    name          VARCHAR2(20),
    diameter      NUMBER);
/
CREATE TYPE nt_sat_t AS TABLE OF satellite_t;
/
CREATE TYPE planet_t AS OBJECT (
    name          VARCHAR2(20),
    mass          NUMBER,
    satellites    nt_sat_t);
/
CREATE TYPE nt_pl_t AS TABLE OF planet_t;
/
CREATE TABLE stars (
    name          VARCHAR2(20),
    age           NUMBER,
    planets       nt_pl_t)
NESTED TABLE planets STORE AS planets_tab
    (NESTED TABLE satellites STORE AS satellites_tab);

```

The preceding example can refer to the inner `satellite` nested table by name because this nested table is a named attribute of an object. However, if the inner nested table is not an attribute, it has no name. The keyword `COLUMN_VALUE` is provided for this case; you use it in place of a name for an inner nested table. For example:

```

CREATE TYPE inner_table AS TABLE OF NUMBER;
/
CREATE TYPE outer_table AS TABLE OF inner_table;
/
CREATE TABLE tab1 (
    col1 NUMBER,
    col2 outer_table)
NESTED TABLE col2 STORE AS col2_ntab
    (NESTED TABLE COLUMN_VALUE STORE AS cv_ntab);

```

Physical attributes for the storage tables can be specified in the nested table clause. For example:

```

CREATE TABLE stars (
    name          VARCHAR2(20),
    age           NUMBER,
    planets       nt_pl_t)
NESTED TABLE planets STORE AS planets_tab

```

```
( PRIMARY KEY (NESTED_TABLE_ID, name)
  ORGANIZATION INDEX COMPRESS
  NESTED TABLE satellites STORE AS satellites_tab );
```

Every nested table storage table contains a column, referenceable by `NESTED_TABLE_ID`, that keys rows in the storage table to the associated row in the parent table. A parent table that is itself a nested table has two system-supplied ID columns: one, referenceable by `NESTED_TABLE_ID`, that keys its rows back to rows in its own parent table, and one hidden column referenced by the `NESTED_TABLE_ID` column in its nested table children.

In the preceding example, nested table `planets` is made an index-organized table (IOT) by adding the `ORGANIZATION INDEX` clause and assigning the nested table a primary key in which the first column is `NESTED_TABLE_ID`. This column contains the ID of the row in the parent table with which a storage table row is associated. Specifying a primary key with `NESTED_TABLE_ID` as the first column and index-organizing the table cause Oracle to physically cluster all the nested table rows that belong to the same parent row, for more efficient access.

Each nested table needs its own table storage clause, so you must have as many nested table storage clauses as you have levels of nested tables in a collection. See ["Nested Table Storage"](#) on page 8-14.

### Assignment and Comparison of Multilevel Collections

As with single-level collections, both the source and the target must be of the same declared data type in assignments of multilevel collections.

Only items whose datatypes are nested table collection types, including multilevel collection types, can be compared. See ["Comparisons of Collections"](#) on page 3-17.

### Constructors for Multilevel Collections

Like single-level collection types, multilevel collection types are created by calling the respective type's constructor method. Like the constructor methods for other user-defined types, a constructor for a multilevel collection type is a system-defined function that has the same name as the type and returns a new instance of it—in this case, a new multilevel collection. Constructor parameters have the names and types of the object type's attributes.

The following example calls the constructor for the multilevel collection type `nt_pl_t`. This type is a nested table of planets, each of which contains a nested table of satellites as an attribute. The constructor for the outer nested table calls the `planet_t` constructor for each planet to be created; each planet constructor calls

the constructor for the satellites nested table type to create its nested table of satellites; and the satellites nested table type constructor calls the `satellite_t` constructor for each satellite instance to be created.

```
INSERT INTO stars
VALUES( 'Sun',23,
      nt_pl_t(
        planet_t(
          'Neptune',
          10,
          nt_sat_t(
            satellite_t('Proteus',67),
            satellite_t('Triton',82)
          )
        ),
        planet_t(
          'Jupiter',
          189,
          nt_sat_t(
            satellite_t('Callisto',97),
            satellite_t('Ganymede', 22)
          )
        )
      )
);
```

## Operations on Collection Datatypes

This section describes the operations on collection datatypes.

### Querying Collections

There are two general ways to query a table that contains a column or attribute of a collection type. One way returns the collections nested in the result rows that contain them. The other way distributes or unnests collections such that each collection element appears on a row by itself.

#### Nesting Results of Collection Queries

The following queries use the `department_persons` table shown in [Example 3-1](#) on page 3-3. The column `dept_emps` is a nested table collection of `person_typ` type. The `dept_emps` collection column appears in the `SELECT` list like an ordinary, scalar column. Querying a collection column in the `SELECT` list like this

nests the elements of the collection in the result row with which the collection is associated.

For example, the following query retrieves the collection of employees. The collection of employees is nested:

```
SELECT d.dept_emps
       FROM department_persons d;
```

```
DEPT_EMPS(IDNO, NAME, PHONE)
```

```
-----
PEOPLE_TYP(PERSON_TYP(1, 'John Smith', '1-800-555-1212'), PERSON_TYP(2, 'Diane Smith',
'1-800-555-1243'))
```

Results are also nested if an object-type column in the `SELECT` list contains a collection attribute, even if that collection is not explicitly listed in the `SELECT` list itself. For example, the query `SELECT * FROM department_persons` would produce a nested result.

### Unnesting Results of Collection Queries

Not all tools or applications are able to deal with results in a nested format. To view Oracle collection data using tools that require a conventional format, you must unnest, or flatten, the collection attribute of a row into one or more relational rows. You can do this by using a `TABLE` expression with the collection. A `TABLE` expression enables you to query a collection in the `FROM` clause like a table. In effect, you join the nested table with the row that contains the nested table.

The `TABLE` expression can be used to query any collection value expression, including transient values such as variables and parameters.

Like the preceding example, the following query retrieves the collection of employees, but the collection is unnested:

```
SELECT e.*
       FROM department_persons d, TABLE(d.dept_emps) e;
```

```
      IDNO NAME                                PHONE
-----
      1 John Smith                            1-800-555-1212
      2 Diane Smith                            1-800-555-1243
```

As the preceding example shows, a `TABLE` expression can have its own table alias. In the example, a table alias for the `TABLE` expression appears in the `SELECT` list to select columns returned by the `TABLE` expression.

The `TABLE` expression uses another table alias to specify the table that contains the collection column that the `TABLE` expression references. The expression `TABLE(d.dept_emps)` specifies the `department_persons` table as containing the `dept_emps` collection column. A `TABLE` expression can use the table alias of any table appearing to the left of it in a `FROM` clause to reference a column of that table. This way of referencing collection columns is called left correlation.

In the example, the `department_persons` table is listed in the `FROM` clause solely to provide a table alias for the `TABLE` expression to use. No columns from the `department_persons` table other than the column referenced by the `TABLE` expression appear in the result

The following example produces rows only for departments that have employees.

```
SELECT d.dept_no, e.*
       FROM department_persons d, TABLE(d.dept_emps) e;
```

To get rows for departments that have no employees, you can use outer-join syntax:

```
SELECT d.dept_no, e.*
       FROM department_persons d, TABLE(d.dept_emps) (+) e;
```

The `(+)` indicates that the dependent join between `department_persons` and `e.dept_emps` should be `NULL`-augmented. That is, there will be rows of `department_persons` in the output for which `e.dept_emps` is `NULL` or empty, with `NULL` values for columns corresponding to `e.dept_emps`.

### Unnesting Queries Containing Table Expression Subqueries

The preceding examples show a `TABLE` expression that contains the name of a collection. Alternatively, a `TABLE` expression can contain a subquery of a collection.

The following example returns the collection of employees whose department number is 101.

```
SELECT *
       FROM TABLE(SELECT d.dept_emps
                   FROM department_persons d
                   WHERE d.dept_no = 101);
```

There are these restrictions on using a subquery in a `TABLE` expression:

- The subquery must return a collection type
- The `SELECT` list of the subquery must contain exactly one item

- The subquery must return only a single collection; it cannot return collections for multiple rows. For example, the subquery `SELECT dept_emps FROM department_persons` succeeds in a `TABLE` expression only if table `department_persons` contains just a single row. If the table contains more than one row, the subquery produces an error.

Here is an example showing a `TABLE` expression used in the `FROM` clause of a `SELECT` embedded in a `CURSOR` expression:

```
SELECT d.dept_no, CURSOR(SELECT * FROM TABLE(d.dept_emps))
FROM department_persons d
WHERE d.dept_no = 101;
```

### Unnesting Queries with Multilevel Collections

Unnesting queries can be used with multilevel collections, too, for both varrays and nested tables. The following example shows an unnesting query on a multilevel nested table collection of nested tables. From a table `stars` in which each star has a nested table of planets and each planet has a nested table of satellites, the query returns the names of all satellites from the inner set of nested tables.

```
SELECT t.name
FROM stars s, TABLE(s.planets) p, TABLE(p.satellites) t;
```

Because no columns of the base table `stars` appear in the `SELECT` list, the query is optimized to run directly against the `satellites` storage table.

Outer-join syntax can also be used with queries of multilevel collections.

**See Also:** ["Viewing Object Data in Relational Form with Unnesting Queries"](#) on page 8-11

## Performing DML Operations on Collections

Oracle supports the following DML operations on nested table columns:

- Inserts and updates that provide a new value for the entire collection
- Piecewise Updates
  - Inserting new elements into the collection
  - Deleting elements from the collection
  - Updating elements of the collection.

Oracle does not support piecewise updates on `VARRAY` columns. However, `VARRAY` columns can be inserted into or updated as an atomic unit.

For piecewise updates of nested table columns, the DML statement identifies the nested table value to be operated on by using the `TABLE` expression.

The following DML statements demonstrate piecewise operations on nested table columns.

```
INSERT INTO TABLE(SELECT d.dept_emps
                    FROM department_persons d
                    WHERE d.dept_no = 101)
VALUES (5, 'Kevin Taylor', '1-800-555-6212');
```

```
UPDATE TABLE(SELECT d.dept_emps
              FROM department_persons d
              WHERE d.dept_no = 101) e
SET VALUE(e) = person_typ(5, 'Kevin Taylor', '1-800-555-6233')
WHERE e.idno = 5;
```

```
DELETE FROM TABLE(SELECT d.dept_emps
                   FROM department_persons d
                   WHERE d.dept_no = 101) e
WHERE e.idno = 5;
```

### Performing DML on Multilevel Collections

For multilevel nested table collections, DML can be done atomically, on the collection as a whole, or piecewise, on selected elements. For multilevel `varray` collections, DML operations can be done only atomically.

**Collections as Atomic Data Items** The section "[Constructors for Multilevel Collections](#)" on page 3-11 shows an example of inserting an entire multilevel collection with an `INSERT` statement. Multilevel collections can also be updated atomically with an `UPDATE` statement. For example, suppose `v_planets` is a variable declared to be of the planets nested table type `nt_pl_t`. The following statement updates `stars` by setting the `planets` collection as a unit to the value of `v_planets`.

```
UPDATE stars s
SET s.planets = :v_planets
WHERE s.name = 'Aurora Borealis';
```

**Piecewise Operations on Nested Tables** Piecewise DML is possible only on nested tables, not on `varrays`.

The following example shows a piecewise insert operation on the `planets` nested table of nested tables: the example inserts a new planet, complete with its own nested table of `satellite_t`:

```
INSERT INTO TABLE( SELECT planets FROM stars WHERE name = 'Sun')
VALUES ('Saturn', 56,
       nt_sat_t(
         satellite_t('Rhea', 83)
       )
);
```

The next example performs a piecewise insert into an inner nested table to add a satellite for a planet. Like the preceding, this example uses a `TABLE` expression containing a subquery that selects the inner nested table to specify the target for the insert.

```
INSERT INTO TABLE( SELECT p.satellites
                    FROM TABLE( SELECT s.planets
                                FROM stars s
                                WHERE s.name = 'Sun') p
                    WHERE p.name = 'Uranus')
VALUES ('Miranda', 31);
```

## Comparisons of Collections

The conditions listed in this section allow comparisons of nested tables. There is no mechanism for comparing varrays. The SQL examples in this section use the nested tables created in [Example 3-2](#) on page 3-5.

### Equal and Not Equal Comparisons

The equal (`=`) and not equal (`<>`) conditions determine whether the input nested tables are identical or not, returning the result as a Boolean value.

Two nested tables are equal if they have the same named type, have the same cardinality, and their elements are equal. Elements are equal depending on whether they are equal by the elements own equality definitions, except for object types which require a map method.

For example:

```
SELECT p.name
       FROM students, TABLE(physics_majors) p
WHERE math_majors = physics_majors;
```

In this example, the nested tables contain `person_typ` objects which have an associated map method. See [Example 1-1](#) on page 1-4.

### In Comparisons

The `IN` condition checks whether a nested table is in a list of nested tables, returning the result as a Boolean value. `NULL` is returned if the nested table is a null nested table.

For example:

```
SELECT p.idno, p.name
   FROM students, TABLE(physics_majors) p
 WHERE physics_majors IN (math_majors, chem_majors);
```

### Subset of Multiset Comparison

The `SUBMULTISET [OF]` condition checks whether a nested table is a subset of another nested table, returning the result as a Boolean value. The `OF` keyword is optional and does not change the functionality of `SUBMULTISET`.

This operator is implemented only for nested tables because this is a multiset function only.

For example:

```
SELECT p.idno, p.name
   FROM students, TABLE(physics_majors) p
 WHERE physics_majors SUBMULTISET OF math_majors;
```

### Member of a Nested Table Comparison

The `MEMBER [OF]` or `NOT MEMBER [OF]` condition tests whether an element is a member of a nested table, returning the result as a Boolean value. The `OF` keyword is optional and has no effect on the output.

For example:

```
SELECT graduation
   FROM students
 WHERE person_typ(12, 'Bob Jones', '1-800-555-1212') MEMBER OF math_majors;
```

where `person_typ (12, 'Bob Jones', '1-800-555-1212')` is an element of the same type as the elements of the nested table `math_majors`.

## Empty Comparison

The `IS [NOT] EMPTY` condition checks whether a given nested table is empty or not empty, regardless of whether any of the elements are `NULL`. If a `NULL` is given for the nested table, the result is `NULL`. The result is returned as a Boolean value.

```
SELECT p.idno, p.name
   FROM students, TABLE(physics_majors) p
 WHERE physics_majors IS NOT EMPTY;
```

## Set Comparison

The `IS [NOT] A SET` condition checks whether a given nested table is composed of unique elements, returning a Boolean value.

For example:

```
SELECT p.idno, p.name
   FROM students, TABLE(physics_majors) p
 WHERE physics_majors IS A SET;
```

## Multisets Operations

This section describes multiset operations with nested tables. The SQL examples in this section use the nested tables created in [Example 3-2](#) on page 3-5.

### See Also:

- *Oracle Database SQL Reference* for information about using operators with nested tables
- "[Comparisons of Objects, REF Variables, and Collections](#)" on page 2-34

## CARDINALITY

The `CARDINALITY` function returns the number of elements in a varray or nested table. The return type is `NUMBER`. If the varray or nested table is a null collection, `NULL` is returned.

For example:

```
SELECT CARDINALITY(math_majors)
   FROM students;
```

## COLLECT

The `COLLECT` function is an aggregate function which would create a multiset from a set of elements. The function would take a column of the element type as input and create a multiset from rows selected. To get the results of this function you must use it within a `CAST` function to specify the output type of `COLLECT`.

### See Also:

- ["CAST"](#) on page 2-35 for an example of the `COLLECT` function
- *Oracle Database SQL Reference* for syntax of the `COLLECT` function

## MULTISET EXCEPT

The `MULTISET EXCEPT` operator inputs two nested tables and returns a nested table whose elements are in the first nested table but not in the second nested table. The input nested tables and the output nested table are all type name equivalent.

The `ALL` or `DISTINCT` options can be used with the operator. The default is `ALL`.

- With the `ALL` option, for `ntab1 MULTISET EXCEPT ALL ntab2`, all elements in `ntab1` other than those in `ntab2` would be part of the result. If a particular element occurs  $m$  times in `ntab1` and  $n$  times in `ntab2`, the result will have  $(m - n)$  occurrences of the element if  $m$  is greater than  $n$  otherwise 0 occurrences of the element.
- With the `DISTINCT` option, any element that is present in `ntab1` which is also present in `ntab2` would be eliminated, irrespective of the number of occurrences.

For example:

```
SELECT math_majors MULTISET EXCEPT physics_majors
   FROM students
 WHERE graduation = '01-JUN-03';
```

## MULTISET INTERSECTION

The `MULTISET INTERSECT` operator returns a nested table whose values are common in the two input nested tables. The input nested tables and the output nested table are all type name equivalent.

There are two options associated with the operator: `ALL` or `DISTINCT`. The default is `ALL`. With the `ALL` option, if a particular value occurs  $m$  times in `ntab1` and  $n$  times in `ntab2`, the result would contain the element  $\text{MIN}(m, n)$  times. With the

**DISTINCT** option the duplicates from the result would be eliminated, including duplicates of **NULL** values if they exist.

For example:

```
SELECT math_majors MULTISSET INTERSECT physics_majors
   FROM students
 WHERE graduation = '01-JUN-03';
```

## MULTISSET UNION

The **MULTISSET UNION** operator returns a nested table whose values are those of the two input nested tables. The input nested tables and the output nested table are all type name equivalent.

There are two options associated with the operator: **ALL** or **DISTINCT**. The default is **ALL**. With the **ALL** option, all elements that are in *ntab1* and *ntab2* would be part of the result, including all copies of **NULLS**. If a particular element occurs *m* times in *ntab1* and *n* times in *ntab2*, the result would contain the element (*m + n*) times. With the **DISTINCT** option the duplicates from the result are eliminated, including duplicates of **NULL** values if they exist.

For example:

```
SELECT math_majors MULTISSET UNION DISTINCT physics_majors
   FROM students
 WHERE graduation = '01-JUN-03';

PEOPLE_TYP(PERSON_TYP(12, 'Bob Jones', '1-800-555-1212'),
           PERSON_TYP(31, 'Sarah Chen', '1-800-555-2212'),
           PERSON_TYP(45, 'Chris Woods', '1-800-555-1213'))

SELECT math_majors MULTISSET UNION ALL physics_majors
   FROM students
 WHERE graduation = '01-JUN-03';

PEOPLE_TYP(PERSON_TYP(12, 'Bob Jones', '1-800-555-1212'),
           PERSON_TYP(31, 'Sarah Chen', '1-800-555-2212'),
           PERSON_TYP(45, 'Chris Woods', '1-800-555-1213'),
           PERSON_TYP(12, 'Bob Jones', '1-800-555-1212'),
           PERSON_TYP(45, 'Chris Woods', '1-800-555-1213'))
```

## POWMULTISET

The `POWMULTISET` function generates all non-empty submultisets from a given multiset. The input to the `POWMULTISET` function could be any expression which evaluates to a multiset. The limit on the cardinality of the multiset argument is 32.

For example:

```
SELECT * FROM TABLE(POWMULTISET( people_typ (
    person_typ(12, 'Bob Jones', '1-800-555-1212'),
    person_typ(31, 'Sarah Chen', '1-800-555-2212'),
    person_typ(45, 'Chris Woods', '1-800-555-1213'))));
```

## POWMULTISET\_BY\_CARDINALITY

The `POWMULTISET_BY_CARDINALITY` function returns all non-empty submultisets of a nested table of the specified cardinality. The output would be rows of nested tables.

`POWMULTISET_BY_CARDINALITY(x, l)` is equivalent to `TABLE(POWMULTISET(x)) p` where `CARDINALITY(value(p)) = l`, where `x` is a multiset and `l` is the specified cardinality.

The first input parameter to the `POWMULTISET_BY_CARDINALITY` could be any expression which evaluates to a nested table. The length parameter should be a positive integer, otherwise an error will be returned. The limit on the cardinality of the nested table argument is 32.

For example:

```
SELECT * FROM TABLE(POWMULTISET_BY_CARDINALITY( people_typ (
    person_typ(12, 'Bob Jones', '1-800-555-1212'),
    person_typ(31, 'Sarah Chen', '1-800-555-2212'),
    person_typ(45, 'Chris Woods', '1-800-555-1213')),2));
```

## SET

The `SET` function converts a nested table into a set by eliminating duplicates, and returns a nested table whose elements are `DISTINCT` from one another. The nested table returned is of the same named type as the input nested table.

For example:

```
SELECT SET(physics_majors)
FROM students
WHERE graduation = '01-JUN-03';
```

---

# Object Support in Oracle Programming Environments

In Oracle you can create object types with SQL data definition language (DDL) commands, and you can manipulate objects with SQL data manipulation language (DML) commands. Object support is built into Oracle application programming environments.

This chapter discusses the following topics:

- [SQL](#)
- [PL/SQL](#)
- [Oracle Call Interface \(OCI\)](#)
- [Pro\\*C/C++](#)
- [Oracle C++ Call Interface \(OCCI\)](#)
- [Oracle Objects For OLE \(OO4O\)](#)
- [Java: JDBC, Oracle SQLJ, JPublisher, and SQLJ Object Types](#)
- [XML](#)

## SQL

Oracle SQL DDL provides the following support for object types:

- Defining object types, nested tables, and arrays
- Specifying privileges
- Specifying table columns of user-defined types
- Creating object tables

Oracle SQL DML provides the following support for object types:

- Querying and updating objects and collections
- Manipulating `REFs`

**See Also:** For a complete description of Oracle SQL syntax, see *Oracle Database SQL Reference*.

## PL/SQL

Object types and subtypes can be used in PL/SQL procedures and functions in most places where built-in types can appear.

The parameters and variables of PL/SQL functions and procedures can be of object types.

You can implement the methods associated with object types in PL/SQL. These methods (functions and procedures) reside on the server as part of a user's schema.

**See Also:** For a complete description of PL/SQL, see the *PL/SQL User's Guide and Reference*.

## Oracle Call Interface (OCI)

OCI is a set of C library functions that applications can use to manipulate data and schemas in an Oracle database. OCI supports both traditional 3GL and object-oriented techniques for database access, as explained in the following sections.

An important component of OCI is a set of calls to manage a workspace called the object cache. The object cache is a memory block on the client side that allows programs to store entire objects and to navigate among them without additional round trips to the server.

The object cache is completely under the control and management of the application programs using it. The Oracle server has no access to it. The application programs using it must maintain data coherency with the server and protect the workspace against simultaneous conflicting access.

OCI provides functions to

- Access objects on the server using SQL.
- Access, manipulate and manage objects in the object cache by traversing pointers or `REFs`.
- Convert Oracle dates, strings and numbers to C data types.
- Manage the size of the object cache's memory.

OCI improves concurrency by allowing individual objects to be locked. It improves performance by supporting complex object retrieval.

OCI developers can use the object type translator to generate the C datatypes corresponding to a Oracle object types.

**See Also:** *Oracle Call Interface Programmer's Guide* for more information about using objects with OCI

## Associative Access in OCI Programs

Traditionally, 3GL programs manipulate data stored in a relational database by executing SQL statements and PL/SQL procedures. Data is usually manipulated on the server without incurring the cost of transporting the data to the client(s). OCI supports this associative access to objects by providing an API for executing SQL statements that manipulate object data. Specifically, OCI enables you to:

- Execute SQL statements that manipulate object data and object type schema information
- Pass object instances, object references (`REFs`), and collections as input variables in SQL statements
- Return object instances, `REFs`, and collections as output of SQL statement fetches
- Describe the properties of SQL statements that return object instances, `REFs`, and collections
- Describe and execute PL/SQL procedures or functions with object parameters or results

- Synchronize object and relational functionality through enhanced commit and rollback functions

**See Also:** ["Associative Access in Pro\\*C/C++"](#) on page 4-7

## Navigational Access in OCI Programs

In the object-oriented programming paradigm, applications model their real-world entities as a set of inter-related objects that form graphs of objects. The relationships between objects are implemented as references. An application processes objects by starting at some initial set of objects, using the references in these initial objects to traverse the remaining objects, and performing computations on each object. OCI provides an API for this style of access to objects, known as navigational access. Specifically, OCI enables you to:

- Cache objects in memory on the client machine
- De-reference an object reference and pin the corresponding object in the object cache. The pinned object is transparently mapped in the host language representation.
- Notify the cache when the pinned object is no longer needed
- Fetch a graph of related objects from the database into the client cache in one call
- Lock objects
- Create, update, and delete objects in the cache
- Flush changes made to objects in the client cache to the database

**See Also:** ["Navigational Access in Pro\\*C/C++"](#) on page 4-7

## Object Cache

To support high-performance navigational access of objects, OCI runtime provides an object cache for caching objects in memory. The object cache supports references (REFs) to database objects in the object cache, the database objects can be identified (that is, pinned) through their references. Applications do not need to allocate or free memory when database objects are loaded into the cache, because the object cache provides transparent and efficient memory management for database objects.

Also, when database objects are loaded into the cache, they are transparently mapped into the host language representation. For example, in the C programming language, the database object is mapped to its corresponding C structure. The object

cache maintains the association between the object copy in the cache and the corresponding database object. Upon transaction commit, changes made to the object copy in the cache are propagated automatically to the database.

The object cache maintains a fast look-up table for mapping `REFs` to objects. When an application de-references a `REF` and the corresponding object is not yet cached in the object cache, the object cache automatically sends a request to the server to fetch the object from the database and load it into the object cache. Subsequent de-references of the same `REF` are faster because they become local cache access and do not incur network round-trips. To notify the object cache that an application is accessing an object in the cache, the application pins the object; when it is finished with the object, it unpins it. The object cache maintains a pin count for each object in the cache. The count is incremented upon a pin call and decremented upon an unpin call. When the pin count goes to zero, it means the object is no longer needed by the application. The object cache uses a least-recently used (LRU) algorithm to manage the size of the cache. When the cache reaches the maximum size, the LRU algorithm frees candidate objects with a pin count of zero.

## Building an OCI Program That Manipulates Objects

When you build an OCI program that manipulates objects, you must complete the following general steps:

1. Define the object types that correspond to the application objects.
2. Execute the SQL DDL statements to populate the database with the necessary object types.
3. Represent the object types in the host language format.

For example, to manipulate instances of the object types in a C program, you must represent these types in the C host language format. You can do this by representing the object types as C structs. You can use a tool provided by Oracle called the Object Type Translator (OTT) to generate the C mapping of the object types. The OTT puts the equivalent C structs in header (\*.h) files. You include these \*.h files in the \*.c files containing the C functions that implement the application.

4. Construct the application executable by compiling and linking the application's \*.c files with the OCI library.

**See Also:** *Oracle Call Interface Programmer's Guide* for tips and techniques for using OCI program effectively with objects

## Defining User-Defined Constructors in C

When defining a user-defined constructor in C, you must specify `SELF` (and you may optionally specify `SELF TDO`) in the `PARAMETERS` clause. On entering the C function, the attributes of the C structure that the object maps to are all initialized to `NULL`. The value returned by the function is mapped to an instance of the user-defined type.

For example:

```
CREATE LIBRARY person_lib TRUSTED AS STATIC
/

CREATE TYPE person AS OBJECT
( name VARCHAR2(30),
  CONSTRUCTOR FUNCTION person(name VARCHAR2) RETURN SELF AS RESULT);
/

CREATE TYPE BODY person IS
  CONSTRUCTOR FUNCTION person(name VARCHAR2) RETURN SELF AS RESULT
  IS EXTERNAL NAME "cons_person_typ" LIBRARY person_lib WITH CONTEXT
  PARAMETERS(context, SELF, name OCIString, name INDICATOR sb4);
END;
/
```

The `SELF` parameter is mapped like an `IN` parameter, so in the case of a `NOT FINAL` type, it is mapped to `(dvoid *)`, not `(dvoid **)`.

The return value's `TDO` must match the `TDO` of `SELF` and is therefore implicit. The return value can never be null, so the return indicator is implicit as well.

## Pro\*C/C++

The Oracle Pro\*C/C++ precompiler allows programmers to use user-defined datatypes in C and C++ programs.

Pro\*C developers can use the Object Type Translator to map Oracle object types and collections into C datatypes to be used in the Pro\*C application.

Pro\*C provides compile time type checking of object types and collections and automatic type conversion from database types to C datatypes.

Pro\*C includes an EXEC SQL syntax to create and destroy objects and offers two ways to access objects in the server:

- SQL statements and PL/SQL functions or procedures embedded in Pro\*C programs.
- An interface to the object cache (described under "[Oracle Call Interface \(OCI\)](#)" on page 4-2), where objects can be accessed by traversing pointers, then modified and updated on the server.

**See Also:** For a complete description of the Pro\*C precompiler, see *Pro\*C/C++ Programmer's Guide*.

## Associative Access in Pro\*C/C++

For background information on associative access, see "[Associative Access in OCI Programs](#)" on page 4-3.

Pro\*C/C++ offers the following capabilities for associative access to objects:

- Support for transient copies of objects allocated in the object cache
- Support for transient copies of objects referenced as input host variables in embedded SQL INSERT, UPDATE, and DELETE statements, or in the WHERE clause of a SELECT statement
- Support for transient copies of objects referenced as output host variables in embedded SQL SELECT and FETCH statements
- Support for ANSI dynamic SQL statements that reference object types through the DESCRIBE statement, to get the object's type and schema information

## Navigational Access in Pro\*C/C++

For background information on navigational access, see "[Navigational Access in OCI Programs](#)" on page 4-4.

Pro\*C/C++ offers the following capabilities to support a more object-oriented interface to objects:

- Support for de-referencing, pinning, and optionally locking an object in the object cache using an embedded SQL OBJECT DEREf statement
- Allowing a Pro\*C/C++ user to inform the object cache when an object has been updated or deleted, or when it is no longer needed, using embedded SQL OBJECT UPDATE, OBJECT DELETE, and OBJECT RELEASE statements
- Support for creating new referenceable objects in the object cache using an embedded SQL OBJECT CREATE statement

- Support for flushing changes made in the object cache to the server with an embedded SQL `OBJECT FLUSH` statement

## Converting Between Oracle Types and C Types

The C representation for objects that is generated by the Oracle Type Translator (OTT) uses OCI types whose internal details are hidden, such as `OCIString` and `OCINumber` for scalar attributes. Collection types and object references are similarly represented using `OCITable`, `OCIArray`, and `OCIRef` types. While using these opaque types insulates you from changes to their internal formats, using such types in a C or C++ application is cumbersome. Pro\*C/C++ provides the following ease-of-use enhancements to simplify use of OCI types in C and C++ applications:

- Object attributes can be retrieved and implicitly converted to C types with the embedded SQL `OBJECT GET` statement.
- Object attributes can be set and converted from C types with the embedded SQL `OBJECT SET` statement.
- Collections can be mapped to a host array with the embedded SQL `COLLECTION GET` statement. Furthermore, if the collection is comprised of scalar types, then the OCI types can be implicitly converted to a compatible C type.
- Host arrays can be used to update the elements of a collection with the embedded SQL `COLLECTION SET` statement. As with the `COLLECTION GET` statement, if the collection is comprised of scalar types, C types are implicitly converted to OCI types.

## Oracle Type Translator (OTT)

The Oracle type translator (OTT) is a program that automatically generates C language structure declarations corresponding to object types. OTT makes it easier to use the Pro\*C precompiler and the OCI server access package.

**See Also:** For complete information about OTT, see *Oracle Call Interface Programmer's Guide* and *Pro\*C/C++ Programmer's Guide*.

## Oracle C++ Call Interface (OCCI)

The Oracle C++ Call Interface (OCCI) is a C++ API that enables you to use the object-oriented features, native classes, and methods of the C++ programming language to access the Oracle database.

The OCCI interface is modeled on the JDBC interface and, like the JDBC interface, is easy to use. OCCI itself is built on top of OCI and provides the power and performance of OCI using an object-oriented paradigm.

OCI is a C API to the Oracle database. It supports the entire Oracle feature set and provides efficient access to both relational and object data, but it can be challenging to use—particularly if you want to work with complex, object datatypes. Object types are not natively supported in C, and simulating them in C is not easy. OCCI addresses this by providing a simpler, object-oriented interface to the functionality of OCI. It does this by defining a set of wrappers for OCI. By working with these higher-level abstractions, developers can avail themselves of the underlying power of OCI to manipulate objects in the server through an object-oriented interface that is significantly easier to program.

The Oracle C++ Call Interface, OCCI, can be roughly divided into three sets of functionalities, namely:

- Associative relational access
- Associative object access
- Navigational access

## OCCI Associative Relational and Object Interfaces

The associative relational API and object classes provide SQL access to the database. Through these interfaces, SQL is executed on the server to create, manipulate, and fetch object or relational data. Applications can access any datatype on the server, including the following:

- Large objects
- Objects/Structured types
- Arrays
- References

## The OCCI Navigational Interface

The navigational interface is a C++ interface that lets you seamlessly access and modify object-relational data in the form of C++ objects without using SQL. The C++ objects are transparently accessed and stored in the database as needed.

With the OCCI navigational interface, you can retrieve an object and navigate through references from that object to other objects. Server objects are materialized as C++ class instances in the application cache.

An application can use OCCI object navigational calls to perform the following functions on the server's objects:

- Create, access, lock, delete, and flush objects
- Get references to the objects and navigate through them

**See Also:** *Oracle C++ Call Interface Programmer's Guide* for a complete account of how to build applications with the Oracle C++ API

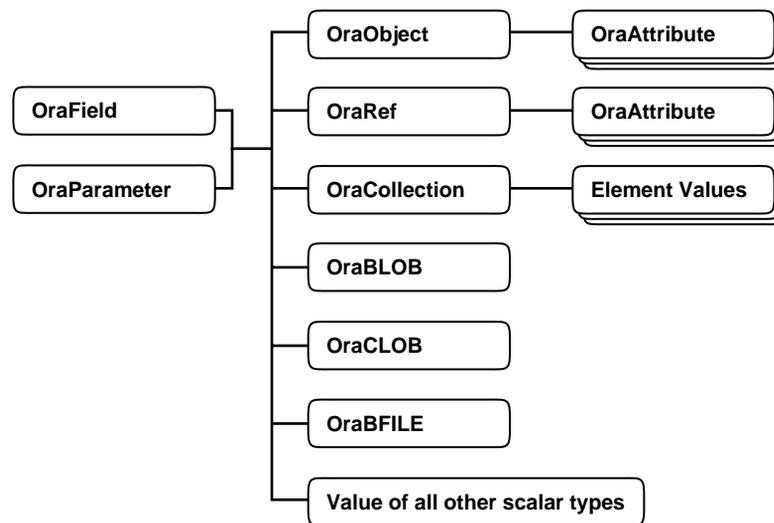
## Oracle Objects For OLE (OO4O)

Oracle Objects for OLE (OO4O) provides full support for accessing and manipulating instances of `REFs`, value instances, variable-length arrays (`VARRAYs`), and nested tables in an Oracle database server.

On Windows systems, you can use Oracle Objects for OLE (OO4O) to write object-oriented database programs in Visual Basic or other environments that support the COM protocol, such as Excel, ActiveX, and Active Server Pages.

**See Also:** The "OO4O Automation Server Reference" section of the Oracle Objects for OLE online help or *Oracle Objects for OLE Developer's Guide* online documentation for detailed information and examples on using OO4O with Oracle objects

[Figure 4-1](#) illustrates the containment hierarchy for value instances of all types in OO4O.

**Figure 4–1 Supported Oracle Datatypes**

Instances of these types can be fetched from the database or passed as input or output variables to SQL statements and PL/SQL blocks, including stored procedures and functions. All instances are mapped to COM Automation Interfaces that provide methods for dynamic attribute access and manipulation. These interfaces may be obtained from:

- The value property of an OraField object in a Dynaset
- The value property of an OraParameter object used as an input or an output parameter in SQL Statements or PL/SQL blocks
- An attribute of an object (REF)
- An element in a collection (varray or a nested table)

## Representing Objects in Visual Basic (OraObject)

The OraObject interface is a representation of an Oracle embedded object or a value instance. It contains a collection interface (OraAttributes) for accessing and manipulating (updating and inserting) individual attributes of a value instance. Individual attributes of an OraAttributes collection interface can be accessed by using a subscript or the name of the attribute.

The following Visual Basic example illustrates how to access attributes of the Address object in the person\_tab table:

```
Dim Address OraObject
Set Person = OraDatabase.CreateDynaset("select * from person_tab", 0&)
Set Address = Person.Fields("Addr").Value
Msgbox Address.Zip
Msgbox.Address.City
```

## Representing REFs in Visual Basic (OraRef)

The OraRef interface represents an Oracle object reference (REF) as well as referenceable objects in client applications. The object attributes are accessed in the same manner as attributes of an object represented by the OraObject interface. OraRef is derived from an OraObject interface by means of the containment mechanism in COM. REF objects are updated and deleted independent of the context they originated from, such as Dynasets. The OraRef interface also encapsulates the functionality for navigating through graphs of objects utilizing the Complex Object Retrieval Capability (COR) in OCI.

**See Also:** *Oracle Call Interface Programmer's Guide* for tips and techniques for using OCI program effectively with objects

## Representing VARRAYs and Nested Tables in Visual Basic (OraCollection)

The OraCollection interface provides methods for accessing and manipulating Oracle collection types, namely variable-length arrays (VARRAYs) and nested tables in OO4O. Elements contained in a collection are accessed by subscripts.

The following Visual Basic example illustrates how to access attributes of the EnameList object from the department table:

```
Dim EnameList OraCollection
Set Person = OraDatabase.CreateDynaset("select * from department", 0&)
Set EnameList = Department.Fields("Enames").Value
'The following loop accesses all elements of the EnameList VArray
For I=1 to I=EnameList.Size
    Msgbox EnameList(I)
Next I
```

## Java: JDBC, Oracle SQLJ, JPublisher, and SQLJ Object Types

Java has emerged as a powerful, modern object-oriented language that provides developers with a simple, efficient, portable, and safe application development platform. Oracle provides two ways to integrate Oracle object features with Java:

JDBC and Oracle SQLJ. These interfaces enable you both to access SQL data from Java and to provide persistent database storage for Java objects.

For an example of using Java APIs with Oracle objects, see the Oracle by Example Series available on the Oracle Technology Network (OTN) Web site.

You can use the followings steps to navigate to the Oracle by Example Series.

- Go to <http://otn.oracle.com>
- Select **Oracle Database** under **Products** from the menu on the left side of the page
- Select **Oracle9i Database Release 1** under **Previous Releases** on the right side of the page
- Under **Technical Information**, select **Oracle9i by Example Series**
- Select **Build Application Components** from the menu on the left side of the page
- Select the **Using Objects to Build an Online Product Catalog** example

You can also search for **Using Objects to Build an Online Product Catalog** on the OTN Web site.

## JDBC Access to Oracle Object Data

JDBC (Java Database Connectivity) is a set of Java interfaces to the Oracle server. Oracle provides tight integration between objects and JDBC. You can map SQL types to Java classes with considerable flexibility.

Oracle JDBC:

- Allows access to objects and collection types (defined in the database) in Java programs through dynamic SQL.
- Translates types defined in the database into Java classes through default or customizable mappings.

Version 2.0 of the JDBC specification supports object-relational constructs such as user-defined (object) types. JDBC materializes Oracle objects as instances of particular Java classes. Using JDBC to access Oracle objects involves creating the Java classes for the Oracle objects and populating these classes. You can either:

- Let JDBC materialize the object as a `STRUCT`. In this case, JDBC creates the classes for the attributes and populates them for you.

- Manually specify the mappings between Oracle objects and Java classes; that is, customize your Java classes for object data. The driver then populates the customized Java classes that you specify, which imposes a set of constraints on the Java classes. To satisfy these constraints, you can choose to define your classes according to either the `SQLData` interface or the `ORAData` interface.

**See Also:** For complete information about JDBC, see the *Oracle Database JDBC Developer's Guide and Reference*.

## SQLJ Access to Oracle Object Data

SQLJ provides access to server objects using SQL statements embedded in the Java code:

- You can use user-defined types in Java programs.
- You can use JPublisher to map Oracle object and collection types into Java classes to be used in the application.
- The object types and collections in the SQL statements are checked at compile time.

**See Also:** For complete information about SQLJ, see the *Oracle Database Java Developer's Guide*.

## Choosing a Data Mapping Strategy

Oracle SQLJ supports either strongly typed or weakly typed Java representations of object types, reference types (`REFs`), and collection types (`varrays` and nested tables) to be used in iterators or host expressions.

Strongly typed representations use a *custom Java class* that corresponds to a particular object type, `REF` type, or collection type and must implement the interface `oracle.sql.ORAData`. The Oracle JPublisher utility can automatically generate such custom Java classes.

Weakly typed representations use the class `oracle.sql.STRUCT` (for objects), `oracle.sql.REF` (for references), or `oracle.sql.ARRAY` (for collections).

## Using JPublisher to Create Java Classes for JDBC and SQLJ Programs

Oracle lets you map Oracle object types, reference types, and collection types to Java classes and preserve all the benefits of strong typing. You can:

- Use JPublisher to automatically generate custom Java classes and use those classes without any change.
- Subclass the classes produced by JPublisher to create your own specialized Java classes.
- Manually code custom Java classes without using JPublisher if the classes meet the requirements stated in the *Oracle Database JPublisher User's Guide*.

We recommend that you use JPublisher and subclass when the generated classes do not do everything you need.

### What JPublisher Produces for a User-Defined Object Type

When you run JPublisher for a user-defined object type, it automatically creates the following:

- A custom object class to act as a type definition to correspond to your Oracle object type  
This class includes getter and setter methods for each attribute. The method names are of the form `getXxx()` and `setXxx()` for attribute `xxx`.  
Also, you can optionally instruct JPublisher to generate wrapper methods in your class that invoke the associated Oracle object methods executing in the server.
- A related custom reference class for object references to your Oracle object type  
This class includes a `getValue()` method that returns an instance of your custom object class, and a `setValue()` method that updates an object value in the database, taking as input an instance of the custom object class.

When you run JPublisher for a user-defined collection type, it automatically creates the following:

- A custom collection class to act as a type definition to correspond to your Oracle collection type  
This class includes overloaded `getArray()` and `setArray()` methods to retrieve or update a collection as a whole, a `getElement()` method and `setElement()` method to retrieve or update individual elements of a collection, and additional utility methods.

JPublisher-produced custom Java classes in any of these categories implement the `ORADData` interface and the `getFactory()` method.

**See Also:** *The Oracle Database JPublisher User's Guide* for more information about using JPublisher.

## Java Object Storage

JPublisher enables you to construct Java classes that map to existing SQL types. You can then access the SQL types from a Java application using JDBC.

You can also go in the other direction. That is, you can create SQL types that map to existing Java classes. This capability enables you to provide persistent storage for Java objects. Such SQL types are called SQL types of Language Java, or SQLJ object types. They can be used as the type of an object, an attribute, a column, or a row in an object table. You can navigationaly access objects of such types—Java objects—through either object references or foreign keys, and you can query and manipulate such objects from SQL.

You create SQLJ types with a `CREATE TYPE` statement as you do other user-defined SQL types. For SQLJ types, two special elements are added to the `CREATE TYPE` statement:

- An `EXTERNAL NAME` phrase, used to identify the Java counterpart for each SQLJ attribute and method and the Java class corresponding to the SQLJ type itself
- A `USING` clause, to specify how the SQLJ type is to be represented to the server. The `USING` clause specifies the interface used to retrieve a SQLJ type and the kind of storage.

For example:

```
CREATE TYPE full_address AS OBJECT (a NUMBER);
/

CREATE OR REPLACE TYPE person_t AS OBJECT
  EXTERNAL NAME 'Person' LANGUAGE JAVA
  USING SQLData (
    ss_no NUMBER (9) EXTERNAL NAME 'socialSecurityNo',
    name varchar(100) EXTERNAL NAME 'name',
    address full_address EXTERNAL NAME 'addrs',
    birth_date date EXTERNAL NAME 'birthDate',
    MEMBER FUNCTION age RETURN NUMBER EXTERNAL NAME 'age () return int',
    MEMBER FUNCTION addressf RETURN full_address
      EXTERNAL NAME 'get_address () return long_address',
    STATIC function createf RETURN person_t EXTERNAL NAME 'create ()
      return Person',
    STATIC function createf (name VARCHAR2, addrs full_address, bDate DATE)
      RETURN person_t EXTERNAL NAME 'create (java.lang.String, Long_address,
```

```
oracle.sql.date) return Person',  
ORDER member FUNCTION compare (in_person person_t) RETURN NUMBER  
EXTERNAL NAME 'isSame (Person) return int')  
/  

```

SQLJ types use the corresponding Java class as the body of the type; you do not specify a type body in SQL to contain implementations of the type's methods as you do with ordinary object types.

### Representing SQLJ Types to the Server

How a SQLJ type is represented to the server and stored depends on the interfaces implemented by the corresponding Java class. Currently, Oracle supports a representation of SQLJ types only for Java classes that implement a `SQLData` or `ORADData` interface. These are represented to the server and are accessible through SQL. A representation for Java classes that implement the `java.io.Serializable` interface is not currently supported.

In a SQL representation, the attributes of the type are stored in columns like attributes of ordinary object types. With this representation, all attributes are public because objects are accessed and manipulated through SQL statements, but you can use triggers and constraints to ensure the consistency of the object data.

For a SQL representation, the `USING` clause must specify either `SQLData` or `ORADData`, and the corresponding Java class must implement one of those interfaces. The `EXTERNAL NAME` clause for attributes is optional.

### Creating SQLJ Object Types

The SQL statements to create SQLJ types and specify their mappings to Java are placed in a file called a **deployment descriptor**. Related SQL constraints and privileges are also specified in this file. The types are created when the file is executed.

Below is an overview of the process of creating SQL versions of Java types/classes:

1. Design the Java types.
2. Generate the Java classes.
3. Create the SQLJ object type statements.
4. Construct the JAR file. This is a single file that contains all the classes needed.
5. Using the `loadjava` utility, install the Java classes defined in the JAR file.
6. Execute the statements to create the SQLJ object types.

### Additional Notes About Mapping

The following are additional notes to consider when mapping of Java classes to SQL types:

- You can map a SQLJ static function to a user-defined constructor in the Java class. The return value of this function is of the user-defined type in which the function is locally defined.
- Java static variables are mapped to SQLJ static methods that return the value of the corresponding static variable identified by `EXTERNAL NAME`. The `EXTERNAL NAME` clause for an attribute is optional with a `SQLData` or `ORADData` representation.
- Every attribute in a SQLJ type of a SQL representation must map to a Java field, but not every Java field must be mapped to a corresponding SQLJ attribute: you can omit Java fields from the mapping.
- You can omit classes: you can map a SQLJ type to a non-root class in a Java class hierarchy without also mapping SQLJ types to the root class and intervening superclasses. Doing this enables you to hide the superclasses while still including attributes and methods inherited from them.

However, you must preserve the structural correspondence between nodes in a class hierarchy and their counterparts in a SQLJ type hierarchy. In other words, for two Java classes `j_A` and `j_B` that are related through inheritance and are mapped to two SQL types `s_A` and `s_B`, respectively, there must be exactly one corresponding node on the inheritance path from `s_A` to `s_B` for each node on the inheritance path from `j_A` to `j_B`.

- You can map a Java class to multiple SQLJ types as long as you do not violate the restriction in the preceding paragraph. In other words, no two SQLJ types mapped to the same Java class can have a common supertype ancestor.
- If all Java classes are not mapped to SQLJ types, it is possible that an attribute of a SQLJ object type might be set to an object of an unmapped Java class. Specifically, to a class occurring above or below the class to which the attribute is mapped in an inheritance hierarchy. If the object's class is a superclass of the attribute's type/class, an error is raised. If it is a subclass of the attribute's type/class, the object is mapped to the most specific type in its hierarchy for which a SQL mapping exists

**See Also:** *The Oracle Database JPublisher User's Guide for JPublisher examples of object mapping*

## Evolving SQLJ Types

The `ALTER TYPE` statement enables you to evolve a type by, for example, adding or dropping attributes or methods.

When a SQLJ type is evolved, an additional validation is performed to check the mapping between the class and the type. If the class and the evolved type match, the type is marked valid. Otherwise, the type is marked as pending validation.

Being marked as pending validation is not the same as being marked invalid. A type that is pending validation can still be manipulated with `ALTER TYPE` and `GRANT` statements, for example.

If a type that has a SQL representation is marked as pending evaluation, you can still access tables of that type using any DML or `SELECT` statement that does not require a method invocation.

You cannot, however, execute DML or `SELECT` statements on tables of a type that has a serializable representation and has been marked as pending validation. Data of a serializable type can be accessed only navigationaly, through method invocations. These are not possible with a type that is pending validation. However, you can still re-evolve the type until it passes validation.

**See Also:** ["Type Evolution"](#) on page 7-7

## Constraints

For SQLJ types having a SQL representation, the same constraints can be defined as for ordinary object types.

Constraints are defined on tables, not on types, and are defined at the column level. The following constraints are supported for SQLJ types having a SQL representation:

- Unique constraints
- Primary Key
- Check constraints
- `NOT NULL` constraints on attributes
- Referential constraints

The `IS OF TYPE` constraint on column substitutability is supported, too, for SQLJ types having a SQL representation.

**See Also:** ["Constraining Substitutability"](#) on page 2-30

## Querying SQLJ Objects

SQLJ types can be queried just like ordinary SQL object types. Methods called in a `SELECT` statement must not attempt to change attribute values.

## Inserting Java Objects

Inserting a row in a table containing a column of a SQLJ type requires a call to the type's constructor function to create a Java object of that type.

The implicit, system-generated constructor can be used, or a static function can be defined that maps to a user-defined constructor in the Java class.

## Updating SQLJ Objects

SQLJ objects can be updated either by using an `UPDATE` statement to modify the value of one or more attributes, or by invoking a method that updates the attributes and returns `SELF`—that is, returns the object itself with the changes made.

For example, suppose that `raise()` is a member function that increments the `salary` field/attribute by a specified amount and returns `SELF`. The following statement gives every employee in the object table `employee_objtab` a raise of 1000:

```
UPDATE employee_objtab SET c=c.raise(1000);
```

A column of a SQLJ type can be set to `NULL` or to another column using the same syntax as for ordinary object types. For example, the following statement assigns column `d` to column `c`:

```
UPDATE employee_reltab SET c=d ;
```

## Defining User-Defined Constructors in Java

When you implement a user-defined constructor in Java, the string supplied as the implementing routine must correspond to a static function. For the return type of the function, specify the Java type mapped to the SQL type.

Here is an example of a type declaration that involves a user-defined constructor implemented in Java:

```
CREATE OR REPLACE TYPE Person1_typ AS OBJECT
  EXTERNAL NAME 'pkg1.J_peson' LANGUAGE JAVA USING SQLData
  ( name VARCHAR2(30),
    age NUMBER,
    CONSTRUCTOR FUNCTION Person1_typ(name VARCHAR2, age NUMBER)
      RETURN Person1_typ AS RESULT
```

```
AS LANGUAGE JAVA
  NAME 'pkg1.J_Person.J_Person(java.lang.String, int) return J_Person'
);
```

## XML

`XMLType` views wrap existing relational and object-relational data in XML formats. These views are similar to object views. Each row of an `XMLType` view corresponds to an `XMLType` instance. The object identifier for uniquely identifying each row in the view can be created using an expression such as `extract()` on the `XMLType` value.

**See Also:** *Oracle XML DB Developer's Guide* for information and examples on using XML with Oracle objects



---

---

# Applying an Object Model to Relational Data

This chapter shows how to write object-oriented applications without changing the underlying structure of your relational data.

The chapter contains these topics:

- [Why Use Object Views](#)
- [Defining Object Views](#)
- [Using Object Views in Applications](#)
- [Nesting Objects in Object Views](#)
- [Identifying Null Objects in Object Views](#)
- [Using Nested Tables and Varrays in Object Views](#)
- [Specifying Object Identifiers for Object Views](#)
- [Creating References to View Objects](#)
- [Modelling Inverse Relationships with Object Views](#)
- [Updating Object Views](#)
- [Applying the Object Model to Remote Tables](#)
- [Defining Complex Relationships in Object Views](#)
- [Object View Hierarchies](#)

## Why Use Object Views

Just as a view is a virtual table, an object view is a virtual object table. Each row in the view is an object: you can call its methods, access its attributes using the dot notation, and create a REF that points to it.

Object views are useful in prototyping or transitioning to object-oriented applications because the data in the view can be taken from relational tables and accessed as if the table were defined as an object table. You can run object-oriented applications without converting existing tables to a different physical structure.

Object views can be used like relational views to present only the data that you want users to see. For example, you might create an object view that presents selected data from an employee table but omits sensitive data about salaries.

Using object views can lead to better performance. Relational data that make up a row of an object view traverse the network as a unit, potentially saving many round trips.

You can fetch relational data into the client-side object cache and map it into C structures or C++ or Java classes, so 3GL applications can manipulate it just like native classes. You can also use object-oriented features like complex object retrieval with relational data.

- By synthesizing objects from relational data, you can query the data in new ways. You can view data from multiple tables by using object de-referencing instead of writing complex joins with multiple tables.
- Because the objects in the view are processed within the server, not on the client, this can result in significantly fewer SQL statements and much less network traffic.
- The object data from object views can be pinned and used in the client side object cache. When you retrieve these synthesized objects in the object cache by means of specialized object-retrieval mechanisms, you reduce network traffic.
- You gain great flexibility when you create an object model within a view in that you can continue to develop the model. If you need to alter an object type, you can simply replace the invalidated views with a new definition.
- Using objects in views does not place any restrictions on the characteristics of the underlying storage mechanisms. By the same token, you are not limited by the restrictions of current technology. For example, you can synthesize objects from relational tables which are parallelized and partitioned.
- You can create different complex data models from the same underlying data.

**See Also:**

- *Oracle Database SQL Reference* for a complete description of SQL syntax and usage.
- *PL/SQL User's Guide and Reference* for a complete discussion of PL/SQL capabilities
- *Oracle Database Java Developer's Guide* for a complete discussion of Java.
- *Oracle Call Interface Programmer's Guide* for a complete discussion of those facilities.

## Defining Object Views

The procedure for defining an object view is:

1. Define an object type, where each attribute of the type corresponds to an existing column in a relational table.
2. Write a query that specifies how to extract the data from relational tables. Specify the columns in the same order as the attributes in the object type.
3. Specify a unique value, based on attributes of the underlying data, to serve as an object identifier, which enables you to create pointers (REFs) to the objects in the view. You can often use an existing primary key.

If you want to be able to update an object view, you may have to take another step, if the attributes of the object type do not correspond exactly to columns in existing tables:

4. Write an `INSTEAD OF` trigger procedure for Oracle to execute whenever an application program tries to update data in the object view. See "[Updating Object Views](#)" on page 5-12.

After these steps, you can use an object view just like an object table.

For example, the following SQL statements define an object view, where each row in the view is an object of type `employee_t`:

```
CREATE TABLE emp_table (  
    empnum    NUMBER (5),  
    ename     VARCHAR2 (20),  
    salary    NUMBER (9,2),  
    job       VARCHAR2 (20));
```

```
CREATE TYPE employee_t AS OBJECT (  
    empno    NUMBER (5),  
    ename    VARCHAR2 (20),  
    salary   NUMBER (9,2),  
    job      VARCHAR2 (20));  
/  
  
CREATE VIEW emp_view1 OF employee_t  
    WITH OBJECT IDENTIFIER (empno) AS  
    SELECT e.empnum, e.ename, e.salary, e.job  
    FROM emp_table e  
    WHERE job = 'Developer';
```

To access the data from the `empnum` column of the relational table, you would access the `empno` attribute of the object type.

## Using Object Views in Applications

Data in the rows of an object view may come from more than one table, but the object still traverses the network in one operation. The instance appears in the client side object cache as a C or C++ structure or as a PL/SQL object variable. You can manipulate it like any other native structure.

You can refer to object views in SQL statements in the same way you refer to an object table. For example, object views can appear in a `SELECT` list, in an `UPDATE-SET` clause, or in a `WHERE` clause.

You can also define object views on object views.

You can access object view data on the client side using the same OCI calls you use for objects from object tables. For example, you can use `OCIObjectPin()` for pinning a `REF` and `OCIObjectFlush()` for flushing an object to the server. When you update or flush to the server an object in an object view, Oracle updates the object view.

**See Also:** See *Oracle Call Interface Programmer's Guide* for more information about OCI calls.

## Nesting Objects in Object Views

An object type can have other object types nested in it as attributes.

If the object type on which an object view is based has an attribute that itself is an object type, then you must provide column objects for this attribute as part of the

process of creating the object view. If column objects of the attribute type already exist in a relational table, you can simply select them; otherwise, you must synthesize the object instances from underlying relational data just as you synthesize the principal object instances of the view. You synthesize, or create, these objects by calling the respective object type's constructor method to create the object instances, and you populate their attributes with data from relational columns that you specify in the constructor.

For example, consider the department table `dept`:

```
CREATE TABLE dept (
  deptno      NUMBER PRIMARY KEY,
  deptname    VARCHAR2(20),
  deptstreet  VARCHAR2(20),
  deptcity    VARCHAR2(10),
  deptstate   CHAR(2),
  deptzip     VARCHAR2(10));
```

You might want to create an object view where the addresses are objects inside the department objects. That would allow you to define reusable methods for address objects, and use them for all kinds of addresses.

**1. Create the type for the address object:**

```
CREATE TYPE address_t AS OBJECT (
  street  VARCHAR2(20),
  city    VARCHAR2(10),
  state   CHAR(2),
  zip     VARCHAR2(10));
/
```

**2. Create the type for the department object:**

```
CREATE TYPE dept_t AS OBJECT (
  deptno    NUMBER,
  deptname  VARCHAR2(20),
  address   address_t );
/
```

**3. Create the view containing the department number, name and address. The address objects are constructed from columns of the relational table.**

```
CREATE VIEW dept_view OF dept_t WITH OBJECT IDENTIFIER (deptno) AS
  SELECT d.deptno, d.deptname,
         address_t(d.deptstreet,d.deptcity,d.deptstate,d.deptzip) AS
         deptaddr
  FROM dept d;
```

## Identifying Null Objects in Object Views

Because the constructor for an object never returns a null, none of the address objects in the preceding view can ever be null, even if the city, street, and so on columns in the relational table are all null. The relational table has no column that specifies whether the department address is null. If we define a convention so that a null `deptstreet` column indicates that the whole address is null, then we can capture the logic using the `DECODE` function, or some other function, to return either a null or the constructed object:

```
CREATE VIEW dept_view AS
  SELECT d.deptno, d.deptname,
         DECODE(d.deptstreet, NULL, NULL,
              address_t(d.deptstreet, d.deptcity, d.deptstate, d.deptzip)) AS deptaddr
  FROM dept d;
```

Using such a technique makes it impossible to directly update the department address through the view, because it does not correspond directly to a column in the relational table. Instead, we would define an `INSTEAD OF` trigger over the view to handle updates to this column.

## Using Nested Tables and Varrays in Object Views

Collections, both nested tables and `VARRAYS`, can be columns in views. You can select these collections from underlying collection columns or you can synthesize them using subqueries. The `CAST-MULTISET` operator provides a way of synthesizing such collections.

### Single-Level Collections in Object Views

Taking the previous example as our starting point, we represent each employee in an `emp` relational table that has the following structure:

```
CREATE TABLE emp (
  empno    NUMBER PRIMARY KEY,
  empname  VARCHAR2(20),
  salary   NUMBER,
  job      VARCHAR2 (20),
  deptno   NUMBER REFERENCES dept(deptno));
```

Using this relational table, we can construct a `dept_view` with the department number, name, address and a collection of employees belonging to the department.

1. Define a nested table type for the employee type `employee_t`:

```
CREATE TYPE employee_list_t AS TABLE OF employee_t;
/
```

2. Define a department type having a department number, name, address, and a nested table of employees:

```
CREATE TYPE dept_t AS OBJECT (
    deptno    NUMBER,
    deptname  VARCHAR2(20),
    address   address_t,
    emp_list  employee_list_t);
/
```

3. Define the object view dept\_view:

```
CREATE VIEW dept_view OF dept_t WITH OBJECT IDENTIFIER (deptno) AS
SELECT d.deptno, d.deptname,
       address_t(d.deptstreet,d.deptcity,d.deptstate,d.deptzip) AS deptaddr,
       CAST( MULTISET (
           SELECT e.empno, e.empname, e.salary, e.job
           FROM emp e
           WHERE e.deptno = d.deptno)
           AS employee_list_t)
       AS emp_list
FROM dept d;
```

The `SELECT` subquery inside the `CAST-MULTISET` block selects the list of employees that belong to the current department. The `MULTISET` keyword indicates that this is a list as opposed to a singleton value. The `CAST` operator casts the result set into the appropriate type, in this case to the `employee_list_t` nested table type.

A query on this view could give us the list of departments, with each department row containing the department number, name, the address object and a collection of employees belonging to the department.

## Multilevel Collections in Object Views

Multilevel collections and single-level collections are created and used in object views in the same way. The only difference is that, for a multilevel collection, you must create an additional level of collections.

The following example builds an object view containing a multilevel collection. The view is based on flat relational tables (that contain no collections). As a preliminary to building the object view, the example creates the object and collection types it uses. An object type (for example, `emp_t`) is defined to correspond to each relational table, with attributes whose types correspond to the types of the

respective table columns. In addition, the employee type has a nested table (attribute) of projects, and the department type has a nested table (attribute) of employees. The latter nested table is a multilevel collection. The `CAST-MULTISET` operator is used in the `CREATE VIEW` statement to build the collections.

These are the underlying relational tables:

```
CREATE TABLE depts
  ( deptno    NUMBER,
    deptname  VARCHAR2(20));

CREATE TABLE emps
  ( ename VARCHAR2(20),
    salary  NUMBER,
    deptname VARCHAR2(20));

CREATE TABLE projects
  ( projname  VARCHAR2(20),
    mgr       VARCHAR2(20));
```

These are the object and collection types the view will use:

```
CREATE TYPE project_t AS OBJECT
  ( projname  VARCHAR2(20),
    mgr       VARCHAR2(20));
/
CREATE TYPE nt_project_t AS TABLE OF project_t;
/
CREATE TYPE emp_t AS OBJECT
  ( ename     VARCHAR2(20),
    salary    NUMBER,
    deptname  VARCHAR2(20),
    projects  nt_project_t );
/
CREATE TYPE nt_emp_t AS TABLE OF emp_t;
/
CREATE TYPE depts_t AS OBJECT
  ( deptno    NUMBER,
    deptname  VARCHAR2(20),
    emps      nt_emp_t );
/
```

The following statement creates the object view:

```
CREATE VIEW v_depts OF depts_t WITH OBJECT IDENTIFIER (deptno) AS
  SELECT d.deptno, d.deptname,
         CAST(MULTISET(SELECT e.ename, e.salary, e.deptname,
```

```

CAST(MULTISET(SELECT p.projname, p.mgr
              FROM projects p
              WHERE p.mgr = e.ename)
     AS nt_project_t)
FROM emps e
WHERE e.deptname = d.deptname)
AS nt_emp_t)
FROM depts d;

```

## Specifying Object Identifiers for Object Views

You can construct pointers (REFs) to the row objects in an object view. Because the view data is not stored persistently, you must specify a set of distinct values to be used as object identifiers. The notion of object identifiers allows the objects in object views to be referenced and pinned in the object cache.

If the view is based on an object table or an object view, then there is already an object identifier associated with each row and you can reuse them. Either omit the `WITH OBJECT IDENTIFIER` clause, or specify `WITH OBJECT IDENTIFIER DEFAULT`.

However, if the row object is synthesized from relational data, you must choose some other set of values.

Oracle lets you specify object identifiers based on the primary key. The set of unique keys that identify the row object is turned into an identifier for the object. These values must be unique within the rows selected out of the view, because duplicates would lead to problems during navigation through object references.

The object view created with the `WITH OBJECT IDENTIFIER` clause has an object identifier derived from the primary key. If the `WITH OBJECT IDENTIFIER DEFAULT` clause is specified, the object identifier is either system generated or primary key based, depending on the underlying table or view definition.

For example, note the definition of the object type `dept_t` and the object view `dept_view` described in ["Single-Level Collections in Object Views"](#) on page 5-6.

Because the underlying relational table has `deptno` as the primary key, each department row has a unique department number. In the view, the `deptno` column becomes the `deptno` attribute of the object type. Once we know that `deptno` is unique within the view objects, we can specify it as the object identifier.

**See Also:** ["Storage Considerations for Object Identifiers \(OIDs\)"](#)  
on page 8-6

## Creating References to View Objects

In the example we have been developing, each object selected out of the `dept_view` view has a unique object identifier derived from the department number value. In the relational case, the foreign key `deptno` in the `emp` employee table matches the `deptno` primary key value in the `dept` department table. We used the primary key value for creating the object identifier in the `dept_view`. This allows us to use the foreign key value in the `emp_view` in creating a reference to the primary key value in `dept_view`.

We accomplish this by using `MAKE_REF` operator to synthesize a primary key object reference. This takes the view or table name to which the reference points and a list of foreign key values to create the object identifier portion of the reference that will match with a particular object in the referenced view.

In order to create an `emp_view` view which has the employee's number, name, salary and a reference to the department in which she works, we need first to create the employee type `emp_t` and then the view based on that type

```
CREATE TYPE emp_t AS OBJECT (  
    empno    NUMBER,  
    ename    VARCHAR2(20),  
    salary   NUMBER,  
    deptref  REF dept_t);  
/  
CREATE OR REPLACE VIEW emp_view OF emp_t WITH OBJECT IDENTIFIER(empno)  
    AS SELECT e.empno, e.empname, e.salary,  
            MAKE_REF(dept_view, e.deptno)  
    FROM emp e;
```

The `deptref` column in the view holds the department reference. We write the following simple query to determine all employees whose department is located in the city of San Francisco:

```
SELECT e.eno, e.salary, e.deptref.deptno  
    FROM emp_view e  
    WHERE e.deptref.address.city = 'San Francisco';
```

Note that we could also have used the `REF` modifier to get the reference to the `dept_view` objects:

```
CREATE OR REPLACE VIEW emp_view OF emp_t WITH OBJECT IDENTIFIER(empno)  
    AS SELECT e.empno, e.empname, e.salary, REF(d)  
    FROM emp e, dept_view d  
    WHERE e.deptno = d.deptno;
```

In this case we join the `dept_view` and the `emp` table on the `deptno` key. The advantage of using `MAKE_REF` operator instead of the `REF` modifier is that in using the former, we can create circular references. For example, we can create employee view to have a reference to the department in which she works, and the department view can have a list of references to the employees who work in that department.

Note that if the object view has a primary key based object identifier, the reference to such a view is primary key based. On the other hand, a reference to a view with system generated object identifier will be a system generated object reference. This difference is only relevant when you create object instances in the OCI object cache and need to get the reference to the newly created objects. This is explained in a later section.

As with synthesized objects, we can also select persistently stored references as view columns and use them seamlessly in queries. However, the object references to view objects cannot be stored persistently.

## Modelling Inverse Relationships with Object Views

Views with objects can be used to model inverse relationships.

### One-to-One Relationships

One-to-one relationships can be modeled with inverse object references. For example, let us say that each employee has a particular computer on her desk, and that the computer belongs to that employee only. A relational model would capture this using foreign keys either from the computer table to the employee table, or in the reverse direction. Using views, we can model the objects so that we have an object reference from the employee to the computer object and also have a reference from the computer object to the employee.

### One-to-Many and Many-to-One Relationships

One-to-many relationships (or many-to-many relationships) can be modeled either by using object references or by embedding the objects. One-to-many relationship can be modeled by having a collection of objects or object references. The many-to-one side of the relationship can be modeled using object references.

Consider the department-employee case. In the underlying relational model, we have the foreign key in the employee table. Using collections in views, we can model the relationship between departments and employees. The department view can have a collection of employees, and the employee view can have a reference to the department (or inline the department values). This gives us both the forward

relation (from employee to department) and the inverse relation (department to list of employees). The department view can also have a collection of references to employee objects instead of embedding the employee objects.

## Updating Object Views

You can update, insert, and delete data in an object view using the same SQL DML you use for object tables. Oracle updates the base tables of the object view if there is no ambiguity.

A view is not directly updatable if its view query contains joins, set operators, aggregate functions, or `GROUP BY` or `DISTINCT` clauses. Also, individual columns of a view are not directly updatable if they are based on pseudocolumns or expression in the view query.

If a view is not directly updatable, you can still update it indirectly using `INSTEAD OF` triggers. To do so, you define an `INSTEAD OF` trigger for each kind of DML statement you want to execute on the view. In the `INSTEAD OF` trigger, you code the operations that must take place on the underlying tables of the view to accomplish the desired change in the view. Then, when you issue a DML statement for which you have defined an `INSTEAD OF` trigger, Oracle transparently runs the associated trigger.

**See Also:** ["Using INSTEAD OF Triggers to Control Mutating and Validation"](#) on page 5-13 for an example of an `INSTEAD OF` trigger

Something you want to be careful of: In an object view hierarchy, `UPDATE` and `DELETE` statements operate polymorphically just as `SELECT` statements do: the set of rows picked out by an `UPDATE` or `DELETE` statement on a view implicitly includes qualifying rows in any subviews of the specified view as well. See ["Object View Hierarchies"](#) on page 5-19 for a discussion of object view hierarchy and examples defining `Student_v` and `Employee_v`.

For example, the following statement, which deletes all persons from `Person_v`, also deletes all students from `Student_v` and all employees from the `Employee_v` view.

```
DELETE FROM Person_v;
```

To exclude subviews and restrict the affected rows just to those in the view actually specified, use the `ONLY` keyword. For example, the following statement updates only persons and not employees or students.

```
UPDATE ONLY(Person_v) SET address = ...
```

## Updating Nested Table Columns in Views

A nested table can be modified by inserting new elements and updating or deleting existing elements. Nested table columns that are virtual or synthesized, as in a view, are not usually updatable. To overcome this, Oracle allows `INSTEAD OF` triggers to be created on these columns.

The `INSTEAD OF` trigger defined on a nested table column (of a view) is fired when the column is modified. Note that if the entire collection is replaced (by an update of the parent row), the `INSTEAD OF` trigger on the nested table column is not fired.

## Using `INSTEAD OF` Triggers to Control Mutating and Validation

`INSTEAD OF` triggers provide a way of updating complex views that otherwise could not be updated. They can also be used to enforce constraints, check privileges and validate the DML. Using these triggers, you can control mutation of the objects created through an object view that might be caused by inserting, updating and deleting.

For instance, suppose we wanted to enforce the condition that the number of employees in a department cannot exceed 10. To enforce this, we can write an `INSTEAD OF` trigger for the employee view. The trigger is not needed for doing the DML because the view can be updated, but we need it to enforce the constraint.

We implement the trigger by means of the following code:

```
CREATE TRIGGER emp_instr INSTEAD OF INSERT ON emp_view
FOR EACH ROW
DECLARE
    dept_var dept_t;
    emp_count integer;
BEGIN
    -- Enforce the constraint
    -- First get the department number from the reference
    UTL_REF.SELECT_OBJECT(:NEW.deptref, dept_var);

    SELECT COUNT(*) INTO emp_count
    FROM emp
    WHERE deptno = dept_var.deptno;
    IF emp_count < 9 THEN
        -- Do the insert
        INSERT INTO emp (empno, empname, salary, deptno)
        VALUES (:NEW.empno, :NEW.ename, :NEW.salary, dept_var.deptno);
    END IF;
END;
```

## Applying the Object Model to Remote Tables

Although you cannot directly access remote tables as object tables, object views let you access remote tables as if they were object tables.

Consider a company with two branches — one in Washington D.C., and another in Chicago. Each site has an employee table. The headquarters in Washington has a department table with the list of all the departments. To get a total view of the entire organization, we can create views over the individual remote tables and then an overall view of the organization.

First, we create an object view for each employee table:

```
CREATE VIEW emp_washington_view (eno,ename,salary)
  AS SELECT e.empno, e.empname, e.salary
     FROM emp@washington_link e;

CREATE VIEW emp_chicago_view
  AS SELECT e.eno, e.name, e.salary
     FROM emp_tab@chicago_link e;
```

We can now create the global view:

```
CREATE VIEW orgnzn_view OF dept_t WITH OBJECT IDENTIFIER (dno)
  AS SELECT d.deptno, d.deptname,
     address_t(d.deptstreet,d.deptcity,d.deptstate,d.deptzip),
     CAST( MULTISET (
           SELECT e.eno, e.ename, e.salary
           FROM emp_washington_view e)
         AS employee_list_t)
     FROM dept d
     WHERE d.deptcity = 'Washington'
 UNION ALL
  SELECT d.deptno, d.deptname,
     address_t(d.deptstreet,d.deptcity,d.deptstate,d.deptzip),
     CAST( MULTISET (
           SELECT e.eno, e.name, e.salary
           FROM emp_chicago_view e)
         AS employee_list_t)
     FROM dept d
     WHERE d.deptcity = 'Chicago';
```

This view has the list of all employees for each department. We use `UNION ALL` because we cannot have two employees working in more than one department.

## Defining Complex Relationships in Object Views

You can define circular references in object views using the `MAKE_REF` operator: `view_A` can refer to `view_B` which in turn can refer to `view_A`. This allows an object view to synthesize a complex structure such as a graph from relational data.

For example, in the case of the department and employee, the department object currently includes a list of employees. To conserve space, we may want to put references to the employee objects inside the department object, instead of materializing all the employees within the department object. We can construct (pin) the references to employee objects, and later follow the references using the dot notation to extract employee information.

Because the employee object already has a reference to the department in which the employee works, an object view over this model contains circular references between the department view and the employee view.

You can create circular references between object views in two different ways:

- **Method 1: Re-create First View After Creating Second View**
  1. Create view A without any reference to view B.
  2. Create view B, which includes a reference to view A.
  3. Replace view A with a new definition that includes the reference to view B.
- **Method 2: Create First View Using FORCE Keyword**
  1. Create view A with the reference to view B using the `FORCE` keyword.
  2. Create view B with reference to view A. When view A is used, it is validated and re-compiled.

Method 2 has fewer steps, but the `FORCE` keyword may hide errors in the view creation. You need to query the `USER_ERRORS` catalog view to see if there were any errors during the view creation. Use this method only if you are sure that there are no errors in the view creation statement.

Also, if errors prevent the views from being recompiled upon use, you must recompile them manually using the `ALTER VIEW COMPILE` command.

We will see the implementation for both the methods.

## Tables and Types to Demonstrate Circular View References

First, we set up some relational tables and associated object types. Although the tables contain some objects, they are not object tables. To access the data objects, we will create object views later.

The emp table stores the employee information:

```
CREATE TABLE emp
( empno    NUMBER PRIMARY KEY,
  empname  VARCHAR2(20),
  salary   NUMBER,
  deptno   NUMBER );
```

The emp\_t type contains a reference to the department. We need a dummy department type so that the emp\_t type creation succeeds.

```
CREATE TYPE dept_t;
/
```

The employee type includes a reference to the department:

```
CREATE TYPE emp_t AS OBJECT
( eno NUMBER,
  ename VARCHAR2(20),
  salary NUMBER,
  deptref REF dept_t );
/
```

We represent the list of references to employees as a nested table:

```
CREATE TYPE employee_list_ref_t AS TABLE OF REF emp_t;
/
```

The department table is a typical relational table:

```
CREATE TABLE dept
( deptno    NUMBER PRIMARY KEY,
  deptname  VARCHAR2(20),
  deptstreet VARCHAR2(20),
  deptcity  VARCHAR2(10),
  deptstate CHAR(2),
  deptzip   VARCHAR2(10) );
```

To create object views, we need object types that map to columns from the relational tables:

```

CREATE TYPE address_t AS OBJECT
( street      VARCHAR2(20),
  city        VARCHAR2(10),
  state       CHAR(2),
  zip         VARCHAR2(10));
/

```

We earlier created an incomplete type; now we fill in its definition:

```

CREATE OR REPLACE TYPE dept_t AS OBJECT
( dno          NUMBER,
  dname        VARCHAR2(20),
  deptaddr     address_t,
  empreflist   employee_list_ref_t);
/

```

## Creating Object Views with Circular References

Now that we have the underlying relational table definitions, we create the object views on top of them.

### Method 1: Re-create First View After Creating Second View

We first create the employee view with a null in the deptref column. Later, we will turn that column into a reference.

```

CREATE VIEW emp_view OF emp_t WITH OBJECT IDENTIFIER(eno)
  AS SELECT e.empno, e.empname, e.salary, NULL
     FROM emp e;

```

Next, we create the department view, which includes references to the employee objects.

```

CREATE VIEW dept_view OF dept_t WITH OBJECT IDENTIFIER(dno)
  AS SELECT d.deptno, d.deptname,
     address_t(d.deptstreet,d.deptcity,d.deptstate,d.deptzip),
     CAST( MULTISET (
           SELECT MAKE_REF(emp_view, e.empno)
           FROM emp e
           WHERE e.deptno = d.deptno)
     AS employee_list_ref_t)
     FROM dept d;

```

We create a list of references to employee objects, instead of including the entire employee object. We now re-create the employee view with the reference to the department view.

```
CREATE OR REPLACE VIEW emp_view OF emp_t WITH OBJECT IDENTIFIER(eno)
  AS SELECT e.empno, e.empname, e.salary,
           MAKE_REF(dept_view, e.deptno)
  FROM emp e;
```

This creates the views.

### Method 2: Create First View Using FORCE Keyword

If we are sure that the view creation statement has no syntax errors, we can use the FORCE keyword to force the creation of the first view without the other view being present.

First, we create an employee view that includes a reference to the department view, which does not exist at this point. This view cannot be queried until the department view is created properly.

```
CREATE OR REPLACE FORCE VIEW emp_view OF emp_t WITH OBJECT IDENTIFIER(eno)
  AS SELECT e.empno, e.empname, e.salary,
           MAKE_REF(dept_view, e.deptno)
  FROM emp e;
```

Next, we create a department view that includes references to the employee objects. We do not have to use the FORCE keyword here, because emp\_view already exists.

```
CREATE OR REPLACE VIEW dept_view OF dept_t WITH OBJECT IDENTIFIER(dno)
  AS SELECT d.deptno, d.deptname,
           address_t(d.deptstreet,d.deptcity,d.deptstate,d.deptzip),
           CAST( MULTISET (
                 SELECT MAKE_REF(emp_view, e.empno)
                 FROM emp e
                 WHERE e.deptno = d.deptno)
           AS employee_list_ref_t)
  FROM dept d;
```

This allows us to query the department view, getting the employee object by de-referencing the employee reference from the nested table empreflist:

```
SELECT Deref(e.COLUMN_VALUE)
  FROM TABLE( SELECT e.empreflist FROM dept_view e WHERE e.dno = 100) e;
```

`COLUMN_VALUE` is a special name that represents the scalar value in a scalar nested table. In this case, `COLUMN_VALUE` denotes the reference to the employee objects in the nested table `empreflist`.

We can also access only the employee number of all those employees whose name begins with `John`.

```
SELECT e.COLUMN_VALUE.eno
       FROM TABLE(SELECT e.empreflist FROM dept_view e WHERE e.dno = 100) e
       WHERE e.COLUMN_VALUE.ename like 'John%';
```

To get a tabular output, unnest the list of references by joining the department table with the items in its nested table:

```
SELECT d.dno, e.COLUMN_VALUE.eno, e.COLUMN_VALUE.ename
       FROM dept_view d, TABLE(d.empreflist) e
       WHERE e.COLUMN_VALUE.ename like 'John%'
       AND d.dno = 100;
```

Finally, we can rewrite the preceding query to use the `emp_view` instead of the `dept_view` to show how you can navigate from one view to the other:

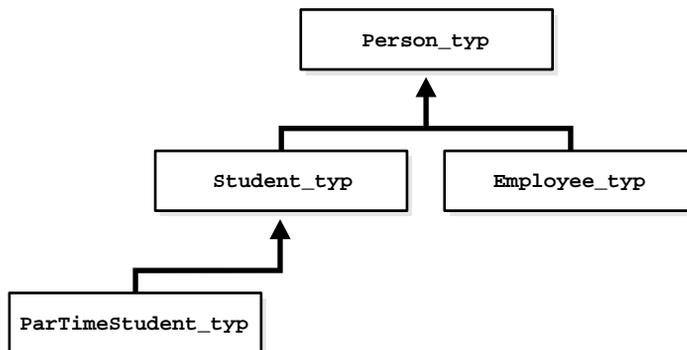
```
SELECT e.deptref.dno, Deref(f.COLUMN_VALUE)
       FROM emp_view e, TABLE(e.deptref.empreflist) f
       WHERE e.deptref.dno = 100
       AND f.COLUMN_VALUE.ename like 'John%';
```

## Object View Hierarchies

An object view hierarchy is a set of object views each of which is based on a different type in a type hierarchy. Subviews in a view hierarchy are created *under* a superview, analogously to the way subtypes in a type hierarchy are created under a supertype.

Each object view in a view hierarchy is populated with objects of a single type, but queries on a given view implicitly address its subviews as well. Thus an object view hierarchy gives you a simple way to frame queries that can return a polymorphic set of objects of a given level of specialization or greater.

For example, suppose you have the following type hierarchy, with `person_typ` as the root:

**Figure 5–1 Object View Hierarchy**

If you have created an object view hierarchy based on this type hierarchy, with an object view built on each type, you can query the object view that corresponds to the level of specialization you are interested in. For instance, you can query the view of `student_typ` to get a result set that contains only students, including part-time students.

You can base the root view of an object view hierarchy on any type in a type hierarchy: you do not need to start the object view hierarchy at the root type. Nor do you need to extend an object view hierarchy to every leaf of a type hierarchy or cover every branch. However, you cannot skip intervening subtypes in the line of descent. Any subview must be based on a direct subtype of the type of its direct superview.

Just as a type can have multiple sibling subtypes, an object view can have multiple sibling subviews. But a subview based on a given type can participate in only one object view hierarchy: two different object view hierarchies cannot each have a subview based on the same subtype.

A subview inherits the object identifier (OID) from its superview. An OID cannot be explicitly specified in any subview.

A root view can explicitly specify an object identifier using the `WITH OBJECT ID` clause. If the OID is system-generated or the clause is not specified in the root view, then subviews can be created only if the root view is based on a table or view that also uses a system generated OID.

The query underlying a view determines whether the view is updatable. For a view to be updatable, its query must contain no joins, set operators, aggregate functions, `GROUP BY`, `DISTINCT`, pseudocolumns, or expressions. The same applies to subviews.

If a view is not updatable, you can define `INSTEAD OF` triggers to perform appropriate DML actions. Note that `INSTEAD OF` triggers are not inherited by subviews.

All views in a view hierarchy must be in the same schema.

---

---

**Note:** You can create views of types that are non-instantiable. A non-instantiable type cannot have instances, so ordinarily there would be no point in creating an object view of such a type. However, a non-instantiable type can have subtypes that *are* instantiable. The ability to create object views of non-instantiable types enables you to base an object view hierarchy on a type hierarchy that contains a non-instantiable type.

---

---

## Creating an Object View Hierarchy

You build an object view hierarchy by creating subviews under a root view. You do this by using the `UNDER` keyword in the `CREATE VIEW` statement.

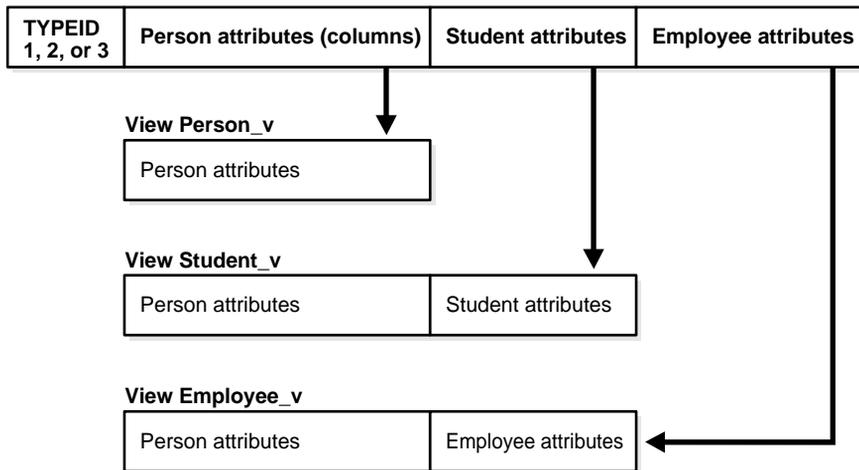
```
CREATE VIEW Student_v OF student_typ UNDER Person_v
AS
SELECT ssn, name, address, deptid, major
FROM AllPersons
WHERE typeid = 2;
```

The same object view hierarchy can be based on different underlying storage models. In other words, a variety of layouts or designs of underlying tables can produce the same object view hierarchy. The design of the underlying storage model has implications for the performance and updatability of the object view hierarchy.

The following examples show three possible storage models. In the first, a flat model, all views in the object view hierarchy are based on the same table. In the second, a horizontal model, each view has a one-to-one correspondence with a different table. And in the third, a vertical model, the views are constructed using joins.

### The Flat Model

In the flat model, all the views in the hierarchy are based on the same table. In the following example, the single table `AllPersons` contains columns for all the attributes of `person_typ`, `student_typ`, or `employee_typ`.

**Figure 5–2 Flat Storage Model for Object View Hierarchy****Table AllPersons**

```
CREATE TABLE AllPersons
( typeid NUMBER(1),
  ssn NUMBER,
  name VARCHAR2(30),
  address VARCHAR2(100),
  deptid NUMBER,
  major VARCHAR2(30),
  empid NUMBER,
  mgr VARCHAR2(30));
```

The `typeid` column identifies the type of each row. Possible values are:

```
1 = person_typ
2 = student_typ
3 = employee_typ
```

```
CREATE TYPE person_typ AS OBJECT
( ssn NUMBER,
  name VARCHAR2(30),
  address VARCHAR2(100)) NOT FINAL;
/
```

```
CREATE TYPE student_typ UNDER person_typ
( deptid NUMBER,
  major VARCHAR2(30)) NOT FINAL;
```

```

/
CREATE TYPE employee_typ UNDER person_typ
( empid NUMBER,
  mgr VARCHAR2(30));
/

```

The following statements create the views that make up the object view hierarchy:

```

CREATE VIEW Person_v OF person_typ
  WITH OBJECT OID(ssn) AS
  SELECT ssn, name, address
  FROM AllPersons
  WHERE typeid = 1;

CREATE VIEW Student_v OF student_typ UNDER Person_v
  AS
  SELECT ssn, name, address, deptid, major
  FROM AllPersons
  WHERE typeid = 2;

CREATE VIEW Employee_v OF employee_typ UNDER Person_v
  AS
  SELECT ssn, name, address, empid, mgr
  FROM AllPersons
  WHERE typeid = 3;

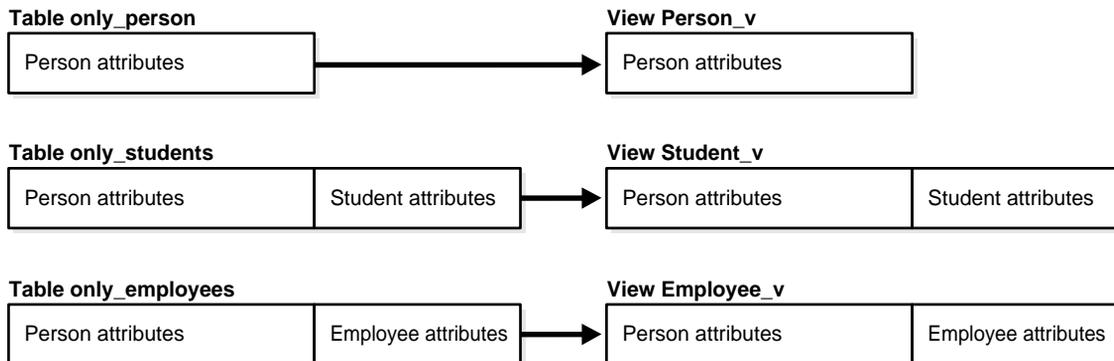
```

The flat model has the advantage of simplicity and poses no obstacles to supporting indexes and constraints. Its drawbacks are:

- A single table cannot contain more than 1000 columns, so the flat model imposes a 1000-column limit on the total number of columns that the object view hierarchy can contain.
- Each row of the table will have NULLs for all the attributes not belonging to its type. Such non-trailing NULLs can adversely affect performance.

### The Horizontal Model

On the horizontal model, each view or subview is based on a different table. In the example, the tables are relational, but they could just as well be object tables for which column substitutability is turned off.

**Figure 5–3 Horizontal Storage Model for Object View Hierarchy**

```
CREATE TABLE only_persons
( ssn NUMBER,
  name VARCHAR2(30),
  address VARCHAR2(100));
```

```
CREATE TABLE only_students
( ssn NUMBER,
  name VARCHAR2(30),
  address VARCHAR2(100),
  deptid NUMBER,
  major VARCHAR2(30));
```

```
CREATE TABLE only_employees
( ssn NUMBER,
  name VARCHAR2(30),
  address VARCHAR2(100),
  empid NUMBER,
  mgr VARCHAR2(30));
```

**These are the views:**

```
CREATE OR REPLACE VIEW Person_v OF person_ttyp
WITH OBJECT OID(ssn) AS
SELECT *
FROM only_persons
```

```
CREATE OR REPLACE VIEW Student_v OF student_ttyp UNDER Person_v
AS
SELECT *
FROM only_students;
```

```
CREATE OR REPLACE VIEW Employee_v OF employee_typ UNDER Person_v
AS
SELECT *
FROM only_employees;
```

The horizontal model is very efficient at processing queries of the form:

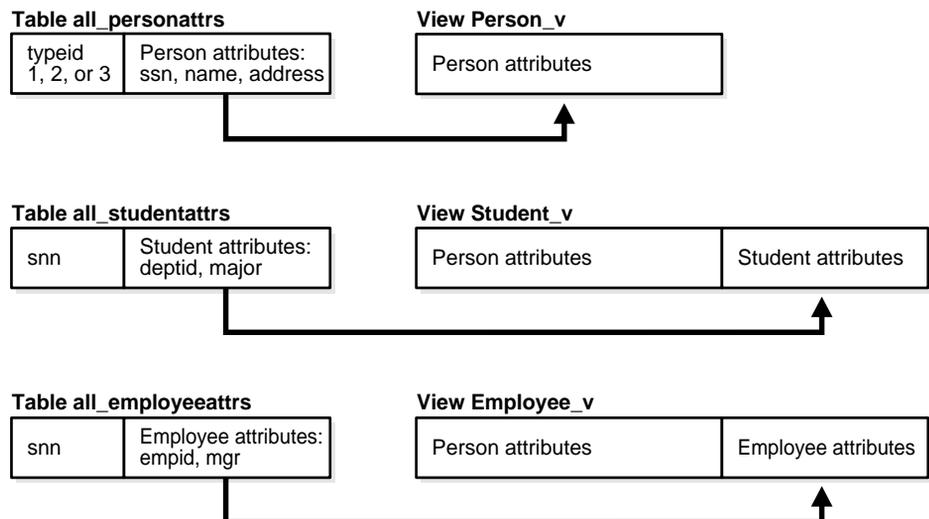
```
SELECT VALUE(p) FROM Person_v p
WHERE VALUE(p) IS OF (ONLY student_typ);
```

Such queries need access only a single physical table to get all the objects of the specific type. The drawbacks of this model are that queries of the sort `SELECT * FROM view` require performing a `UNION` over all the underlying tables and projecting the rows over just the columns in the specified view. (See ["Querying a View in a Hierarchy"](#) on page 5-26.) Also, indexes on attributes (and unique constraints) must span multiple tables, and support for this does not currently exist.

## The Vertical Model

In the vertical model, there is a physical table corresponding to each view in the hierarchy, but each physical table stores only those attributes that are unique to its corresponding subtype.

**Figure 5-4 Vertical Storage Model for Object View Hierarchy**



```
CREATE TABLE all_personattrs
( typeid NUMBER,
  ssn NUMBER,
  name VARCHAR2(30),
  address VARCHAR2(100));

CREATE TABLE all_studentattrs
( ssn NUMBER,
  deptid NUMBER,
  major VARCHAR2(30));

CREATE TABLE all_employeeattrs
( ssn NUMBER,
  empid NUMBER,
  mgr VARCHAR2(30));

CREATE OR REPLACE VIEW Person_v OF person_typ
WITH OBJECT OID(ssn) AS
  SELECT ssn, name, address
  FROM all_personattrs
  WHERE typeid = 1;

CREATE OR REPLACE VIEW Student_v OF student_typ UNDER Person_v
AS
  SELECT x.ssn, x.name, x.address, y.deptid, y.major
  FROM all_personattrs x, all_studentattrs y
  WHERE x.typeid = 2 AND x.ssn = y.ssn;

CREATE OR REPLACE VIEW Employee_v OF employee_typ UNDER Person_v
AS
  SELECT x.ssn, x.name, x.address, y.empid, y.mgr
  FROM all_personattrs x, all_employeeattrs y
  WHERE x.typeid = 3 AND x.ssn = y.ssn;
```

The vertical model can efficiently process queries of the kind `SELECT * FROM root_view`, and it is possible to index individual attributes and impose unique constraints on them. However, to re-create an instance of a type, a join over OIDs must be performed for each level that the type is removed from the root in the hierarchy.

### Querying a View in a Hierarchy

You can query any view or subview in an object view hierarchy; rows are returned for the declared type of the view that you query and for any of that type's subtypes.

So, for instance, in an object view hierarchy based on the `person_ttyp` type hierarchy, you can query the view of `person_ttyp` to get a result set that contains all persons, including students and employees; or you can query the view of `student_ttyp` to get a result set that contains only students, including part-time students.

In the `SELECT` list of a query, you can include either functions such as `REF()` and `VALUE()` that return an object instance, or you can specify object attributes of the view's declared type, such as the `name` and `ssn` attributes of `person_ttyp`.

If you specify functions, to return object instances, the query returns a polymorphic result set: that is, it returns instances of both the view's declared type and any subtypes of that type.

For example, the following query returns instances of persons, employees, and students of all types, as well as `REFs` to those instances.

```
SELECT REF(p), VALUE(p) FROM Person_v p;
```

If you specify individual attributes of the view's declared type in the `SELECT` list or do a `SELECT *`, again the query returns rows for the view's declared type and any subtypes of that type, but these rows are projected over columns for the attributes of the view's declared type, and only those columns are used. In other words, the subtypes are represented only with respect to the attributes they inherit from and share with the view's declared type.

So, for example, the following query returns rows for all persons and rows for employees and students of all types, but the result uses only the columns for the attributes of `person_ttyp`—namely, `name`, `ssn`, and `address`. It does not show rows for attributes added in the subtypes, such as the `deptid` attribute of `student_ttyp`.

```
SELECT * FROM Person_v;
```

To exclude subviews from the result, use the `ONLY` keyword. The `ONLY` keyword confines the selection to the declared type of the view that you are querying:

```
SELECT VALUE(p) FROM ONLY(Person_v) p;
```

## Privileges for Operations on View Hierarchies

Generally, a query on a view with subviews requires only the `SELECT` privilege on the view being referenced and does not require any explicit privileges on subviews. For example, the following query requires only `SELECT` privileges on `Person_v` but not on any of its subviews.

```
SELECT * FROM Person_v;
```

However, a query that selects for any attributes added in subtypes but not used by the root type requires the `SELECT` privilege on all subviews as well. Such subtype attributes may hold sensitive information that should reasonably require additional privileges to access.

The following query, for example, requires `SELECT` privileges on `Person_v` and also on `Student_v`, `Employee_v` (and on any other subview of `Person_v`) because the query selects object instances and thus gets all the attributes of the subtypes.

```
SELECT VALUE(p) FROM Person_v p;
```

To simplify the process of granting `SELECT` privileges on an entire view hierarchy, you can use the `HIERARCHY` option. Specifying the `HIERARCHY` option when granting a user `SELECT` privileges on a view implicitly grants `SELECT` privileges on all current and future subviews of the view as well. For example:

```
GRANT SELECT ON Person_v TO scott WITH HIERARCHY OPTION;
```

A query that excludes rows belonging to subviews also requires `SELECT` privileges on all subviews. The reason is that information about which rows belong exclusively to the most specific type of an instance may be sensitive, so the system requires `SELECT` privileges on subviews for queries (such as the following one) that exclude all rows from subviews.

```
SELECT * FROM ONLY(Person_v);
```

---

# Managing Oracle Objects

This chapter explains how Oracle objects work in combination with the rest of the database, and how to perform DML and DDL operations on them. It contains the following major sections:

- [Privileges on Object Types and Their Methods](#)
- [Dependencies and Incomplete Types](#)
- [Synonyms for User-Defined Types](#)
- [Performance Tuning](#)
- [Tools Providing Support for Objects](#)
- [Utilities Providing Support for Objects](#)

## Privileges on Object Types and Their Methods

Privileges for object types exist at the system level and the schema object level.

### System Privileges for Object Types

Oracle defines the following system privileges for object types:

- `CREATE TYPE` enables you to create object types in your own schema
- `CREATE ANY TYPE` enables you to create object types in any schema
- `ALTER ANY TYPE` enables you to alter object types in any schema
- `DROP ANY TYPE` enables you to drop named types in any schema
- `EXECUTE ANY TYPE` enables you to use and reference named types in any schema
- `UNDER ANY TYPE` enables you to create subtypes under any non-final object types
- `UNDER ANY VIEW` enables you to create subviews under any object view

The `CONNECT` and `RESOURCE` roles include the `CREATE TYPE` system privilege. The `DBA` role includes all of these privileges.

### Schema Object Privileges

Two schema object privileges apply to object types:

- `EXECUTE` on an object type enables you to use the type to:
  - Define a table.
  - Define a column in a relational table.
  - Declare a variable or parameter of the named type.

`EXECUTE` lets you invoke the type's methods, including the constructor.

Method execution and the associated permissions are the same as for stored PL/SQL procedures.

- `UNDER` enables you to create a subtype or subview under the type or view on which the privilege is granted

The `UNDER` privilege on a subtype or subview can be granted only if the grantor has the `UNDER` privilege on the direct supertype or superview `WITH GRANT OPTION`.

The phrase `WITH HIERARCHY OPTION` grants a specified object privilege on all subobjects of the object. This option is meaningful only with the `SELECT` object privilege granted on an object view in an object view hierarchy. In this case, the privilege applies to all subviews of the view on which the privilege is granted.

## Using Types in New Types or Tables

In addition to the permissions detailed in the previous sections, you need specific privileges to:

- Create types or tables that use types created by other users.
- Grant use of your new types or tables to other users.

You must have the `EXECUTE ANY TYPE` system privilege, or you must have the `EXECUTE` object privilege for any type you use in defining a new type or table. You must have received these privileges explicitly, not through roles.

If you intend to grant access to your new type or table to other users, you must have either the required `EXECUTE` object privileges with the `GRANT` option or the `EXECUTE ANY TYPE` system privilege with the option `WITH ADMIN OPTION`. You must have received these privileges explicitly, not through roles.

## Example: Privileges on Object Types

Assume that three users exist with the `CONNECT` and `RESOURCE` roles: `USER1`, `USER2`, and `USER3`.

`USER1` performs the following DDL in the `USER1` schema:

```
CREATE TYPE type1 AS OBJECT ( attr1 NUMBER );
/
CREATE TYPE type2 AS OBJECT ( attr2 NUMBER );
/
GRANT EXECUTE ON type1 TO user2;
GRANT EXECUTE ON type2 TO user2 WITH GRANT OPTION;
```

`USER2` performs the following DDL in the `USER2` schema:

```
CREATE TABLE tab1 OF user1.type1;
CREATE TYPE type3 AS OBJECT ( attr3 user1.type2 );
/
CREATE TABLE tab2 (col1 user1.type2 );
```

The following statements succeed because `USER2` has `EXECUTE` on `USER1`'s `TYPE2` with the `GRANT` option:

```
GRANT EXECUTE ON type3 TO user3;
GRANT SELECT ON tab2 TO user3;
```

However, the following grant fails because `USER2` does not have `EXECUTE` on `USER1.TYPE1` with the `GRANT` option:

```
GRANT SELECT ON tab1 TO user3;
```

`USER3` can successfully perform the following actions:

```
CREATE TYPE type4 AS OBJECT (attr4 user2.type3);
/
CREATE TABLE tab3 OF type4;
```

## Privileges on Type Access and Object Access

While object types only make use of `EXECUTE` privilege, object tables use all the same privileges as relational tables:

- `SELECT` lets you access an object and its attributes from the table.
- `UPDATE` lets you modify attributes of objects in the table.
- `INSERT` lets you add new objects to the table.
- `DELETE` lets you delete objects from the table.

Similar table and column privileges regulate the use of table columns of object types.

Selecting columns of an object table does not require privileges on the type of the object table. Selecting the entire row object, however, does.

Consider the following schema:

```
CREATE TYPE emp_type as object (
    eno      NUMBER,
    ename    VARCHAR2(36));
/
CREATE TABLE emp OF emp_type;
```

and the following two queries:

```
SELECT VALUE(e) FROM emp e;
SELECT eno, ename FROM emp;
```

For either query, Oracle checks the user's `SELECT` privilege for the `emp` table. For the first query, the user needs to obtain the `emp_type` type information to interpret

the data. When the query accesses the `emp_type` type, Oracle checks the user's `EXECUTE` privilege.

Execution of the second query, however, does not involve named types, so Oracle does not check type privileges.

Additionally, using the schema from the previous section, `USER3` can perform the following queries:

```
SELECT t.col1.attr2 from user2.tab2 t;
```

```
SELECT t.attr4.attr3.attr2 FROM tab3 t;
```

Note that in both selects by `USER3`, `USER3` does not have explicit privileges on the underlying types, but the statement succeeds because the type and table owners have the necessary privileges with the `GRANT` option.

Oracle checks privileges on the following requests, and returns an error if the requestor does not have the privilege for the action:

- Pinning an object in the object cache using its `REF` value causes Oracle to check `SELECT` privilege on the object table containing the object and `EXECUTE` privilege on the object type.

**See Also:** *Oracle Call Interface Programmer's Guide* for tips and techniques for using OCI program effectively with objects

- Modifying an existing object or flushing an object from the object cache, causes Oracle to check `UPDATE` privilege on the destination object table. Flushing a new object causes Oracle to check `INSERT` privilege on the destination object table.
- Deleting an object causes Oracle to check `DELETE` privilege on the destination table.
- Invoking a method causes Oracle to check `EXECUTE` privilege on the corresponding object type.

Oracle does not provide column level privileges for object tables.

## Dependencies and Incomplete Types

Types can depend upon each other for their definitions. For example, you might want to define object types `employee` and `department` in such a way that one

attribute of `employee` is the department the employee belongs to and one attribute of `department` is the employee who manages the department.

Types that depend on each other in this way, either directly or through intermediate types, are called mutually dependent. In a diagram that uses arrows to show the dependency relationships among a set of types, connections among mutually dependent types form a loop.

To define such a circular dependency, you must use `REFs` for at least one segment of the circle.

For example, you can define the following types:

```
CREATE TYPE department;
/

CREATE TYPE employee AS OBJECT (
    name    VARCHAR2(30),
    dept    REF department,
    supv    REF employee );
/

CREATE TYPE emp_list AS TABLE OF employee;
/

CREATE TYPE department AS OBJECT (
    name    VARCHAR2(30),
    mgr     REF employee,
    staff   emp_list );
/
```

This is a legal set of mutually dependent types and a legal sequence of SQL DDL statements. Oracle compiles it without errors.

Notice that the preceding code creates the type `department` twice. The first statement:

```
CREATE TYPE department;
/
```

is an optional, incomplete declaration of `department` that serves as a placeholder for the `REF` attribute of `employee` to point to. The declaration is incomplete in that it omits the `AS OBJECT` phrase and lists no attributes or methods. These are specified later in the full declaration that completes the type. In the meantime, `department` is created as an incomplete object type. This enables the compilation of `employee` to proceed without errors.

To complete an incomplete type, you execute a `CREATE TYPE` statement that specifies the attributes and methods of the type, as shown at the end of the example. Complete an incomplete type after all the types that it refers to are created.

If you do not create incomplete types as placeholders, types that refer to the missing types still compile, but the compilation proceeds with errors.

For example, if `department` did not exist at all, Oracle would create it as an incomplete type and compile `employee` with errors. Then `employee` would be recompiled the next time that some operation attempts to access it. This time, if all the types it depends on are created and its dependencies are satisfied, it will compile without errors.

Incomplete types also enable you to create types that contain `REF` attributes to a subtype that has not yet been created. To create such a supertype, first create an incomplete type of the subtype to be referenced. Create the complete subtype after you create the supertype.

A subtype is just a specialized version of its direct supertype and consequently has an explicit dependency on it. To ensure that subtypes are not left behind after a supertype is dropped, all subtypes must be dropped first: a supertype cannot be dropped until all its subtypes are dropped.

## Completing Incomplete Types

When all the types that an incomplete type refers to have been created, there is no longer any need for the incomplete type to remain incomplete, and you should complete the declaration of the type. Completing the type recompiles it and enables the system to release various locks.

You must complete an incomplete object type as an object type: you cannot complete an object type as a collection type (a nested table type or an array type). The only alternative to completing a type declaration is to drop the type.

You must also complete any incomplete types that Oracle creates for you because you did not explicitly create them yourself. The example in the preceding section explicitly creates `department` as an incomplete type. If `department` were not explicitly created as an incomplete type, Oracle would create it as one so that the `employee` type can compile (with errors). You must complete the declaration of `department` as an object type whether you or Oracle declared it as an incomplete type.

## Manually Recompiling a Type

If a type was created with compilation errors, and you attempt an operation on it, such as creating tables or inserting rows, you may receive an error. You need to recompile type `typename` before attempting the operation. To manually recompile a type, execute an `ALTER TYPE typename COMPILE` statement. After you have successfully compiled the type, attempt the operation again.

## Type Dependencies of Substitutable Tables and Columns

A substitutable table or column of type `T` is dependent not only on `T` but on all subtypes of `T` as well. This is because a hidden column is added to the table for each attribute added in a subtype of `T`. The hidden columns are added even if the substitutable table or column contains no data of that subtype.

So, for example, a persons table of type `person_typ` is dependent not only on `person_typ` but also on the `person_typ` subtypes `student_typ` and `part_time_student_typ`.

If you attempt to drop a subtype that has a dependent type, table, or column, the `DROP TYPE` statement returns an error and aborts. For example, trying to drop `part_time_student_typ` will raise an error because of the dependent persons table.

If dependent tables or columns exist but contain no data of the type that you want to drop, you can use the `VALIDATE` keyword to drop the type. The `VALIDATE` keyword causes Oracle to check for actual stored instances of the specified type and to drop the type if none are found. Hidden columns associated with attributes unique to the type are removed as well.

For example, the first `DROP TYPE` statement in the following example fails because `part_time_student_typ` has a dependent table (`persons`). But if `persons` contains no instances of `part_time_student_typ` (and no other dependent table or column does, either), the `VALIDATE` keyword causes the second `DROP TYPE` statement to succeed:

```
-- Following generates an error due to presence of Persons table
DROP TYPE part_time_student_typ;
-- Following succeeds if there are no stored instances of part_time_student_typ
DROP TYPE part_time_student_typ VALIDATE;
```

---

---

**Note:** Oracle recommends that you always use the `VALIDATE` option while dropping subtypes.

---

---

### The FORCE Option

The `DROP TYPE` statement also has a `FORCE` option that causes the type to be dropped even though it may have dependent types or tables. The `FORCE` option should be used only with great care, as any dependent types or tables that do exist are marked invalid and become inaccessible when the type is dropped. Data in a table that is marked invalid because a type it depends on has been dropped can never be accessed again. The only action that can be performed on such a table is to drop it.

**See Also:** ["Type Evolution"](#) on page 7-7 for information about how to alter a type

## Synonyms for User-Defined Types

Just as you can create synonyms for tables, views, and various other schema objects, you can also define synonyms for user-defined types.

Synonyms for types have the same advantages as synonyms for other kinds of schema objects: they provide a location-independent way to reference the underlying schema object. An application that uses public type synonyms can be deployed without alteration in any schema of a database without having to qualify a type name with the name of the schema in which the type was defined.

**See Also:** *Oracle Database Administrator's Guide* for more information on synonyms in general

## Creating a Type Synonym

You create a type synonym with a `CREATE SYNONYM` statement. For example, these statements create a type `typ1` and then create a synonym for it:

```
CREATE TYPE typ1 AS OBJECT (x number);  
/  
CREATE SYNONYM syn1 FOR typ1;
```

Synonyms can be created for collection types, too. The following example creates a synonym for a nested table type:

```
CREATE TYPE typ2 AS TABLE OF NUMBER;  
/  
CREATE SYNONYM syn2 FOR typ2;
```

You create a public synonym by using the `PUBLIC` keyword:

```
CREATE TYPE shape AS OBJECT ( name VARCHAR2(10) );  
/  
CREATE PUBLIC SYNONYM pub_shape FOR shape;
```

The `REPLACE` option enables you to have the synonym point to a different underlying type. For example, the following statement causes `syn1` to point to type `typ2` instead of the type it formerly pointed to:

```
CREATE OR REPLACE SYNONYM syn1 FOR typ2;
```

## Using a Type Synonym

You can use a type synonym anywhere that you can refer to a type. For instance, you can use a type synonym in a DDL statement to name the type of a table column or type attribute. In the following example, synonym `syn1` is used to specify the type of an attribute in type `typ3`:

```
CREATE TYPE typ1 AS OBJECT (x number);  
/  
CREATE SYNONYM syn1 FOR typ1;  
  
CREATE TYPE typ3 AS OBJECT ( a syn1 );  
/
```

The next example shows a type synonym `syn1` used to call the constructor of the user-defined type `typ1`, for which `syn1` is a synonym. The statement returns an object instance of `typ1`:

```
SELECT syn1(0) FROM dual;
```

In the following example, `syn2` is a synonym for a nested table type. The example shows the synonym used in place of the actual type name in a `CAST` expression:

```
SELECT CAST(MULTISET(SELECT eno FROM USER3.EMP) AS syn2) FROM dual;
```

Type synonyms can be used in the following kinds of statements:

- **DML statements:** `SELECT`, `INSERT`, `UPDATE`, `DELETE`, `FLASHBACK TABLE`, `EXPLAIN PLAN`, and `LOCK TABLE`
- **DDL statements:** `AUDIT`, `NOAUDIT`, `GRANT`, `REVOKE`, and `COMMENT`

## Describing Schema Objects That Use Synonyms

If a type or table has been created using type synonyms, the `DESCRIBE` command will show the synonyms in place of the types they represent. Similarly, catalog

views, such as `USER_TYPE_ATTRS`, that show type names will show the associated type synonym names in their place.

You can query the catalog view `USER_SYNONYMS` to find out the underlying type of a type synonym.

### Dependents of Type Synonyms

A type that directly or indirectly references a synonym in its type declaration is a dependent of that synonym. Thus, in the following example, type `typ3` is a dependent type of synonym `syn1`.

```
CREATE TYPE typ3 AS OBJECT ( a syn1 );
```

Other kinds of schema objects that reference synonyms in their DDL statements also become dependents of those synonyms. An object that depends on a type synonym depends on both the synonym and on the synonym's underlying type.

A synonym's dependency relationships affect your ability to drop or rename the synonym. Dependent schema objects are also affected by some operations on synonyms. The following sections describe these various ramifications.

### Restriction on Replacing a Type Synonym

You can replace a synonym only if it has no dependent tables or valid user defined types. Replacing a synonym is equivalent to dropping it and then re-creating a new synonym with the same name.

### Dropping Type Synonyms

You drop a synonym with the `DROP SYNONYM` statement. For example:

```
CREATE SYNONYM syn4 FOR typ1;
```

```
DROP SYNONYM syn4;
```

You cannot drop a type synonym if it has table or valid user-defined types as dependents unless you use the `FORCE` option. The `FORCE` option causes any columns that directly or indirectly depend on the synonym to be marked unused, just as if the actual types of the columns were dropped. (A column indirectly depends on a synonym if, for instance, the synonym is used to specify the type of an attribute of the declared type of the column.)

Any dependent schema objects of a dropped synonym are invalidated. They can be revalidated by creating a local object of the same name as the dropped synonym or by creating a new public synonym with same name.

Dropping the underlying base type of a type synonym has the same effect on dependent objects as dropping the synonym.

### Renaming Type Synonyms

You can rename a type synonym with the `RENAME` statement. Renaming a synonym is equivalent to dropping it and then re-creating it with a new name. You cannot rename a type synonym if it has dependent tables or valid user-defined types. The following example fails because synonym `syn1` has a dependent user-defined type:

```
RENAME syn1 TO syn3;
```

### Public Type Synonyms and Local Schema Objects

You cannot create a local schema object that has the same name as a public synonym if the public synonym has a dependent table or valid user-defined type in the local schema where you want to create the new schema object. Nor can you create a local schema object that has the same name as a private synonym in the same schema.

For instance, in the following example, table `shape_tab` is a dependent table of public synonym `pub_shape` because the table has a column that uses the synonym in its type definition. Consequently, the attempt to create a table that has the same name as public synonym `pub_shape`, in the same schema as the dependent table, fails:

```
-- Following uses public synonym pub_shape
CREATE TABLE shape_tab ( c1 pub_shape );
-- Following is not allowed
CREATE TABLE pub_shape ( c1 NUMBER );
```

## Performance Tuning

When tuning objects, the following items need to be addressed:

- How objects and object views consume CPU and memory resources during runtime
- How to monitor memory and CPU resources during runtime
- How to manage large numbers of objects

Some of the key performance factors are the following:

- `DBMS_STATS` package to collect statistics
- `tkprof` to profile execution of SQL commands
- `EXPLAIN PLAN` to generate the query plans

**See Also:** *Oracle Database Performance Tuning Guide* for details on measuring and tuning the performance of your application

## Tools Providing Support for Objects

This section describes several Oracle tools that provide support for Oracle objects.

### JDeveloper

JDeveloper is a full-featured, integrated development environment for creating multitier Java applications. It enables you to develop, debug, and deploy Java client applications, dynamic HTML applications, web and application server components and database stored procedures based on industry-standard models.

JDeveloper provides powerful features in the following areas:

- Oracle Business Components for Java
- Web Application Development
- Java Client Application Development
- Java in the Database
- Component-Based Development with JavaBeans
- Simplified Database Access
- Visual Integrated Development Environment
- Java Language Support

JDeveloper runs on Windows platforms. It provides a standard GUI based Java development environment that is well integrated with Oracle Application Server and Database.

### Business Components for Java (BC4J)

Supporting standard EJB and CORBA deployment architectures, Oracle Business Components for Java simplifies the development, delivery, and customization of Java business applications for the enterprise. Oracle Business Components for Java

is an application component framework providing developers a set of reusable software building blocks that manage all the common facilities required to:

- Author and test business logic in components which integrate with relational databases
- Reuse business logic through multiple SQL-based views of data
- Access and update the views from servlets, JavaServer Pages (JSPs), and thin-Java Swing clients
- Customize application functionality in layers without requiring modification of the delivered application

### **JPublisher**

JPublisher is a utility, written entirely in Java, that generates Java classes to represent the following user-defined database entities in your Java program:

- Database object types
- Database reference (REF) types
- Database collection types (varrays or nested tables)
- PL/SQL packages

JPublisher enables you to specify and customize the mapping of database object types, reference types, and collection types (varrays or nested tables) to Java classes, in a strongly typed paradigm.

**See Also:** *Oracle Database JPublisher User's Guide*

## **Utilities Providing Support for Objects**

This section describes several Oracle utilities that provide support for Oracle objects.

### **Import/Export of Object Types**

The Export and Import utilities move data into and out of Oracle databases. They also back up or archive data and aid migration to different releases of the Oracle RDBMS.

Export and Import support object types. Export writes object type definitions and all of the associated data to the dump file. Import then re-creates these items from the dump file.

**Types** The definition statements for derived types are exported. On an Import, a subtype may be created before the supertype definition has been imported. In this case, the subtype will be created with compilation errors, which may be ignored. The type will be revalidated after its supertype is created.

**Object View Hierarchies** View definitions for all views belonging to a view hierarchy are exported

### SQL\*Loader

The SQL\*Loader utility moves data from external files into tables in an Oracle database. The files may contain data consisting of basic scalar datatypes, such as INTEGER, CHAR, or DATE, as well as complex user-defined datatypes such as row and column objects (including objects that have object, collection, or REF attributes), collections, and LOBs. Currently, SQL\*Loader supports single-level collections only: you cannot yet use SQL\*Loader to load multilevel collections, that is, collections whose elements are, or contain, other collections.

SQL\*Loader uses control files, which contain SQL\*Loader data definition language (DDL) statements, to describe the format, content, and location of the datafiles.

SQL\*Loader provides two approaches to loading data:

- **Conventional path loading**, which uses the SQL INSERT statement and a bind array buffer to load data into database tables
- **Direct path loading**, which uses the Direct Path Load API to write data blocks directly to the database on behalf of the SQL\*Loader client.

Direct path loading does not use a SQL interface and thus avoids the overhead of processing the associated SQL statements. Consequently, direct path loading tends to provide much better performance than conventional path loading.

Either approach can be used to load data of supported object and collection datatypes.

**See Also:** *Oracle Database Utilities* for instructions on how to use SQL\*Loader



---

# Advanced Topics for Oracle Objects

The other chapters in this book discuss the topics that you need to get started with Oracle objects. The topics in this chapter are of interest once you start applying object-relational techniques to large-scale applications or complex schemas.

The chapter contains these topics:

- [Storage of Objects](#)
- [Creating Indexes on Typeids or Attributes](#)
- [Type Evolution](#)
- [Transient and Generic Types](#)
- [User-Defined Aggregate Functions](#)
- [Partitioning Tables That Contain Oracle Objects](#)

## Storage of Objects

Oracle automatically maps the complex structure of object types into the simple rectangular structure of tables.

### Leaf-Level Attributes

An object type is like a tree structure, where the branches represent the attributes. Attributes that are objects sprout subbranches for their own attributes.

Ultimately, each branch ends at an attribute that is a built-in type; such as `NUMBER`, `VARCHAR2`, or `REF`; or a collection type, such as `VARRAY` or nested table. Each of these leaf-level attributes of the original object type is stored in a table column.

The leaf-level attributes that are not collection types are called the leaf-level scalar attributes of the object type.

### How Row Objects Are Split Across Columns

In an object table, Oracle stores the data for every leaf-level scalar or `REF` attribute in a separate column. Each `VARRAY` is also stored in a column, unless it is too large. Oracle stores leaf-level attributes of nested table types in separate tables associated with the object table. You must declare these tables as part of the object table declaration. See "[Internal Layout of VARRAYs](#)" on page 7-4 and "[Internal Layout of Nested Tables](#)" on page 7-4.

When you retrieve or change attributes of objects in an object table, Oracle performs the corresponding operations on the columns of the table. Accessing the value of the object itself produces a copy of the object, by invoking the default constructor for the type, using the columns of the object table as arguments.

Oracle stores the system-generated object identifier in a hidden column. Oracle uses the object identifier to construct `REFs` to the object.

### Hidden Columns for Tables with Column Objects

When a table is defined with a column of an object type, Oracle adds hidden columns to the table for the object type's leaf-level attributes. Each object-type column also has a corresponding hidden column to store the `NULL` information for the column objects (that is, the atomic nulls of the top-level and the nested objects).

## Hidden Columns for Substitutable Columns and Tables

A substitutable column or object table has a hidden column not only for each attribute of the column's object type but also for each attribute added in any subtype of the object type. These columns store the values of those attributes for any subtype instances inserted in the substitutable column.

For example, a substitutable column of `person_typ` will have associated with it a hidden column for each of the attributes of `person_typ`: `idno`, `name`, and `phone`. It will also have hidden columns for attributes of the subtypes of `person_typ`. For example, the attributes `dept_id` and `major` (for `student_typ`) and `number_hours` (for `part_time_student_typ`).

When a subtype is created, hidden columns for attributes added in the subtype are automatically added to tables containing a substitutable column of any of the new subtype's ancestor types. These retrofit the tables to store data of the new type. If, for some reason, the columns cannot be added, creation of the subtype is rolled back.

When a subtype is dropped with the `VALIDATE` option to `DROP TYPE`, all such hidden columns for attributes unique to the subtype are automatically dropped as well if they do not contain data.

A substitutable column also has associated with it a hidden type discriminant column. This column contains an identifier, called a `typeid`, that identifies the most specific type of each object in the substitutable column. Typically, a `typeid` (`RAW`) is one byte, though it can be as big as four bytes for a large hierarchy.

You can find the `typeid` of a specified object instance using the function `SYS_TYPEID`. For example, suppose that the substitutable object table `person_obj_table` contains three rows, as shown in [Example 2-1](#) on page 2-25.

The following query retrieves `typeids` of object instances stored in the table:

```
SELECT name, SYS_TYPEID(VALUE(p)) typeid
   FROM person_obj_table p;
```

NAME	TYPEID
Bob Jones	01
Joe Lane	02
Kim Patel	03

The catalog views `USER_TYPES`, `DBA_TYPES` and `ALL_TYPES` contain a `TYPEID` column (not hidden) that gives the `typeid` value for each type. You can join on this

column to get the type names corresponding to the typeids in a type discriminant column.

**See Also:** "[SYS\\_TYPEID](#)" on page 2-38 for more information about `SYS_TYPEID` and typeids

## REFs

When Oracle constructs a `REF` to a row object, the constructed `REF` is made up of the object identifier (OID), some metadata of the object table, and, optionally, the `ROWID`.

The size of a `REF` in a column of `REF` type depends on the storage properties associated with the column. For example, if the column is declared as a `REF WITH ROWID`, Oracle stores the `ROWID` in the `REF` column. The `ROWID` hint is ignored for object references in constrained `REF` columns.

If column is declared as a `REF` with a `SCOPE` clause, the column is made smaller by omitting the object table metadata and the `ROWID`. A scoped `REF` is 16 bytes long.

If the object identifier is primary-key based, Oracle may create one or more internal columns to store the values of the primary key depending on how many columns comprise the primary key.

---

---

**Note:** When a `REF` column references row objects whose object identifiers are derived from primary keys, we refer to it as a primary-key-based `REF` or `pkREF`. Columns containing `pkREFs` must be scoped or have a referential constraint.

---

---

## Internal Layout of Nested Tables

The rows of a nested table are stored in a separate storage table. Each nested table column has a single associated storage table, not one for each row. The storage table holds all the elements for all of the nested tables in that column. The storage table has a hidden `NESTED_TABLE_ID` column with a system-generated value that lets Oracle map the nested table elements back to the appropriate row.

You can speed up queries that retrieve entire collections by making the storage table index-organized. Include the `ORGANIZATION INDEX` clause inside the `STORE AS` clause.

A nested table type can contain objects or scalars:

- If the elements are objects, the storage table is like an object table: the top-level attributes of the object type become the columns of the storage table. But because a nested table row has no object identifier column, you cannot construct `REFs` to objects in a nested table.
- If the elements are scalars, the storage table contains a single column called `COLUMN_VALUE` that contains the scalar values.

**See Also:** ["Nested Table Storage"](#) on page 8-14

## Internal Layout of VARRAYs

All the elements of a `VARRAY` are stored in a single column. Depending upon the size of the array, it may be stored inline or in a `BLOB`. See [Storage Considerations for Varrays](#) on page 8-13 for details.

## Creating Indexes on Typeids or Attributes

This section discusses the use of indexes on typeids and attributes.

### Indexing a Type Discriminant Column

Using the `SYS_TYPEID` function, you can build an index on the hidden type discriminant column that every substitutable column has. The type discriminant column contains typeids that identify the most specific type of every object instance stored in the substitutable column. This information is used by the system to evaluate queries that use the `IS OF` predicate to filter by type, but you can access the typeids for your own purposes using the `SYS_TYPEID` function.

---



---

**Note:** Generally, a type discriminant column contains only a small number of distinct typeids: at most, there can be only as many as there are types in the related type hierarchy. The low cardinality of this column makes it a good candidate for a bitmap index.

---



---

For example, the following statement creates a bitmap index on the type discriminant column underlying the substitutable `contact` column of table `contacts`. Function `SYS_TYPEID` is used to reference the type discriminant column:

```
CREATE BITMAP INDEX typeid_idx ON contacts (SYS_TYPEID(contact));
```

## Indexing Subtype Attributes of a Substitutable Column

You can build an index on attributes of any of the types that can be stored in a substitutable column. Attributes of subtypes can be referenced in the `CREATE INDEX` statement by using the `TREAT` function to filter out types other than the desired subtype (and its subtypes); you then use the dot notation to specify the desired attribute.

For example, the following statement creates an index on the `major` attribute of all students in the `contacts` table. The declared type of the `contact` column is `person_typ`, of which `student_typ` is a subtype, so the column may contain instances of `person_typ`, `student_typ`, and subtypes of either one:

```
CREATE INDEX major1_idx ON contacts
  (TREAT(contact AS student_typ).major);
```

`student_typ` is the type that first defined the `major` attribute: the `person_typ` supertype does not have it. Consequently, all the values in the hidden column for the `major` attribute are values for `student_typ` or `parttimestudent_typ` authors (a `student_typ` subtype). This means that the hidden column's values are identical to the values returned by the `TREAT` expression, which returns `major` values for all students, including student subtypes: both the hidden column and the `TREAT` expression list majors for students and nulls for authors of other types. The system exploits this fact and creates index `major1_idx` as an ordinary B-tree index on the hidden column.

Values in a hidden column are identical to the values returned by a `TREAT` expression like the preceding one only if the type named as the target of the `TREAT` function (`student_typ`) is the type that first defined the attribute. If the target of the `TREAT` function is a subtype that merely inherited the attribute, as in the following example, the `TREAT` expression will return non-null `major` values for the subtype (part-time students) but not for its supertype (other students).

```
CREATE INDEX major2_idx ON contacts
  (TREAT(contact AS part_time_student_typ).major);
```

Here the values stored in the hidden column for `major` may be different from the results of the `TREAT` expression. Consequently, an ordinary B-tree index cannot be created on the underlying column. In a case like this, Oracle treats the `TREAT` expression like any other function-based expression and tries to create the index as a function-based index on the result.

The following example, like the previous one, creates a function-based index on the `major` attribute of part-time students, but in this case the hidden column for `major` is associated with a substitutable object table `person_obj_table`:

```
CREATE INDEX major3_idx ON person_obj_table p
  (TREAT(VALUE(p) AS part_time_student_typ).major);
```

## Type Evolution

Changing a user-defined type is called type evolution. You can make the following changes to a user-defined type:

- Add and drop attributes
- Add and drop methods
- Modify a numeric attribute to increase its length, precision, or scale
- Modify a varying length character attribute to increase its length
- Change a type's `FINAL` and `INSTANTIABLE` properties
- Modify limit and size of `VARRAYS`
- Modify length, precision, and scale of collection elements

Changes to a type affect things that reference the type. For example, if you add a new attribute to a type, data in a column of that type must be presented so as to include the new attribute.

Schema objects that directly or indirectly reference a type and are affected by a change to it are called dependents of the type. A type can have these kinds of dependents:

- Table
- Type or subtype
- Program unit (PL/SQL block): procedure, function, package, trigger
- Indextype
- View (including object view)
- Function-based index
- Operator

How a dependent schema object is affected by a change to a type depends on the dependent object and on the nature of the change to the type.

All dependent program units, views, operators and indextypes are marked invalid when a type is modified. The next time one of these invalid schema objects is referenced, it is revalidated using the new type definition. If the object recompiles

successfully, it becomes valid and can be used again. Depending on the change to the type, function-based indexes may be dropped or disabled and need to be rebuilt.

If a type has dependent tables, then, for each attribute added to a type, one or more internal columns are added to the table depending on the new attribute's type. New attributes are added with `NULL` values. For each dropped attribute, the columns associated with that attribute are dropped. For each modified attribute, the length, precision, or scale of its associated column is changed accordingly.

These changes mainly involve updating the tables' metadata (information about a table's structure, describing its columns and their types) and can be done quickly. However, the data in those tables must be updated to the format of the new type version as well. Updating this data can be time-consuming if there is a lot of it, so the `ALTER TYPE` command has options to let you choose whether to convert all dependent table data immediately or to leave it in the old format to be converted piecemeal as it is updated in the course of business.

The `CASCADE` option for `ALTER TYPE` propagates a type change to dependent types and tables. See "[ALTER TYPE Statement for Type Evolution](#)" on page 7-15. `CASCADE` itself has options that let you choose whether to convert table data to the new type format as part of the propagation: the option `INCLUDING TABLE DATA` converts the data; the option `NOT INCLUDING TABLE DATA` does not convert it. By default, the `CASCADE` option converts the data. In any case, table data is always returned in the format of the latest type version. If the table data is stored in the format of an earlier type version, Oracle converts the data to the format of the latest version before returning it, even though the format in which the data is actually stored is not changed until the data is rewritten.

You can retrieve the definition of the latest type from the system view `USER_SOURCE`. You can view definitions of all versions of a type in the `USER_TYPE_VERSIONS` view.

The following example changes `person_typ` by adding one attribute and dropping another. The `CASCADE` keyword propagates the type change to dependent types and tables, but the phrase `NOT INCLUDING TABLE DATA` prevents conversion of the related data.

```
CREATE TABLE person_obj_table OF person_typ;

INSERT INTO person_obj_table
VALUES (person_typ(12, 'Bob Jones', '111-555-1212'));

SELECT value(p) FROM person_obj_table p;
```

```

VALUE(P)(IDNO, NAME, PHONE)
-----
PERSON_TYP(12, 'Bob Jones', '111-555-1212')

ALTER TYPE person_typ
  ADD ATTRIBUTE (email VARCHAR2(80)),
  DROP ATTRIBUTE phone CASCADE NOT INCLUDING TABLE DATA;

-- disconnect and reconnect to accommodate the type change

-- The data of table person_obj_table has not been converted yet, but
-- when the data is retrieved, Oracle returns the data based on
-- the latest type version. The new attribute is initialized to NULL.
SELECT value(p) FROM person_obj_table p;

VALUE(P)(IDNO, NAME, EMAIL)
-----
PERSON_TYP(12, 'Bob Jones', NULL)

```

During `SELECT` statements, even though column data may be converted to the latest type version, the converted data is not written back to the column. If a certain user-defined type column in a table is retrieved often, you should consider converting that data to the latest type version to eliminate redundant data conversions. Converting is especially beneficial if the column contains a `VARRAY` attribute because a `VARRAY` typically takes more time to convert than an object or nested table column.

You can convert a column of data by issuing an `UPDATE` statement to set the column to itself. For example:

```
UPDATE dept_tab SET emp_array_col = emp_array_col;
```

You can convert all columns in a table by using `ALTER TABLE UPGRADE DATA`.

For example:

```
ALTER TYPE person_typ ADD ATTRIBUTE (photo BLOB)
  CASCADE NOT INCLUDING TABLE DATA;
ALTER TABLE dept_tab UPGRADE INCLUDING DATA;
```

## Changes Involved When a Type Is Altered

Only structural changes to a type affect dependent data and require the data to be converted. Changes that are confined to a type's method definitions or behavior (in the type body, where the type's methods are implemented) do not.

These possible changes to a type are structural:

- Adding an attribute
- Dropping an attribute
- Modifying the length, precision, or scale of an attribute
- Changing the finality of a type (which determines whether subtypes can be derived from it) from `FINAL` to `NOT FINAL` or from `NOT FINAL` to `FINAL`.

These changes result in new versions of the altered type and all its dependent types and require the system to add, drop, or modify internal columns of dependent tables as part of the process of converting to the new version.

When you make any of these kinds of changes to a type that has dependent types or tables, the effects of propagating the change are not confined only to metadata but affect data storage arrangements and require the data to be converted.

Besides converting data, you may also need to make other changes. For example, if a new attribute is added to a type, and the type body invokes the type's constructor, then each constructor in the type body must be modified to specify a value for the new attribute. Similarly, if a new method is added, then the type body must be replaced to add the implementation of the new method. The type body can be modified by using the `CREATE OR REPLACE TYPE BODY` statement.

## Steps to Change a Type

This section describes the steps required to make a change to a type. Assume we have the following schema based on the `person_type` type:

```
CREATE TYPE people_typ AS TABLE OF person_typ;
/
CREATE TYPE department_typ AS OBJECT (
    manager    person_typ,
    employee   people_typ);
/
CREATE TABLE department OF department_typ
    NESTED TABLE employee STORE AS employee_store_nt;
```

1. Issue an `ALTER TYPE` statement to alter the type.

The default behavior of an `ALTER TYPE` statement without any option specified is to check if there is any object dependent on the target type. The statement aborts if any dependent object exists. Optional keywords allow cascading the type change to dependent types and tables.

In the following code, conversion of table data is deferred by adding the phrase NOT INCLUDING TABLE DATA.

```
-- Add new attributes to person_typ and propagate the change
-- to employee_store_nt and department_typ
ALTER TYPE person_typ ADD ATTRIBUTE (photo BLOB, email VARCHAR2(80))
    CASCADE NOT INCLUDING TABLE DATA;
```

2. Use CREATE OR REPLACE TYPE BODY to update the corresponding type body to make it current with the new type definition.
3. Upgrade dependent tables to the latest type version and convert the tables' data.

```
ALTER TABLE department UPGRADE INCLUDING DATA;
```

4. Alter dependent PL/SQL program units as needed to take account of changes to the type.
5. Use OTT or JPUB (or another tool) to generate new header files for applications, depending on whether the application is written in C or Java.

Adding a new attribute to a supertype also increases the number of attributes in all its subtypes because these inherit the new attribute. Inherited attributes always precede declared (locally defined) attributes, so adding a new attribute to a supertype causes the ordinal position of all declared attributes of any subtype to be incremented by one recursively. The mappings of the altered type must be updated to include the new attributes. OTT and JPUB do this. If you use some other tool, you must be sure that the type headers are properly synchronized with the type definition in the server; otherwise, unpredictable behavior may result.

6. Modify application code as needed and rebuild the application.

## Validating a Type

When the system executes an ALTER TYPE statement, it first validates the requested type change syntactically and semantically to make sure it is legal. The system performs the same validations as for a CREATE TYPE statement plus some additional ones. For example, it checks to be sure an attribute being dropped is not used as a partitioning key. If the new spec of the target type or any of its dependent types fails the type validations, the ALTER TYPE statement aborts. No new type version is created, and all dependent objects remain unchanged.

If dependent tables exist, further checking is done to ensure that restrictions relating to the tables and any indexes are observed. Again, if the `ALTER TYPE` statement fails the check of table-related restrictions, then the type change is aborted, and no new version of the type is created.

When multiple attributes are added in a single `ALTER TYPE` statement, they are added in the order specified. Multiple type changes can be specified in the same `ALTER TYPE` statement, but no attribute name or method signature can be specified more than once in the statement. For example, adding and modifying the same attribute in a single statement is not allowed.

For example:

```
CREATE TYPE mytype AS OBJECT (attr1 NUMBER, attr2 NUMBER);
/
ALTER TYPE mytype ADD ATTRIBUTE (attr3 NUMBER),
      DROP ATTRIBUTE attr2,
      ADD ATTRIBUTE attr4 NUMBER CASCADE;
```

The resulting definition for `mytype` becomes:

```
(attr1 NUMBER, attr3 NUMBER, attr4 NUMBER);
```

The following `ALTER TYPE` statement, which attempts to make multiple changes to the same attribute (`attr5`), is invalid:

```
-- invalid ALTER TYPE statement
ALTER TYPE mytype ADD ATTRIBUTE (attr5 NUMBER, attr6 CHAR(10)),
      DROP ATTRIBUTE attr5;
```

The following are other notes on validation constraints, table restrictions, and assorted information about the various kinds of changes that can be made to a type.

### Dropping an Attribute

- Dropping all attributes from a root type is not allowed. You must instead drop the type. Because a subtype inherits all the attributes from its supertype, dropping all the attributes from a subtype does not reduce its attribute count to zero; thus, dropping all attributes declared locally in a subtype is allowed.
- Only an attribute declared locally in the target type can be dropped. You cannot drop an inherited attribute from a subtype. Instead, drop the attribute from the type where it is locally declared.
- Dropping an attribute which is part of a table partitioning or sub-partitioning key in a table is not allowed.

- Dropping an attribute of a primary key OID of an object table or an index-organized table (IOT) is not allowed.
- When an attribute is dropped, the column corresponding to the dropped attribute is dropped.
- Indexes, statistics, constraints, and any referential integrity constraints referencing a dropped attribute are removed.

### **Modifying Attribute Type to Increase the Length, Precision, or Scale**

- Expanding the length of an attribute referenced in a function-based index, clustered key or domain index on a dependent table is not allowed.

### **Dropping a Method**

- You can drop a method only from the type in which the method is defined (or redefined): You cannot drop an inherited method from a subtype, and you cannot drop an redefined method from a supertype.
- If a method is not redefined, dropping it using the `CASCADE` option removes the method from the target type and all subtypes. However, if a method is redefined in a subtype, the `CASCADE` will fail and roll back. For the `CASCADE` to succeed, you must first drop each redefined method from the subtype that defines it and only then drop the method from the supertype.

You can consult the `USER_DEPENDENCIES` table to find all the schema objects, including types, that depend on a given type. You can also run the `DBMS_UTILITY.GET_DEPENDENCY` utility to find the dependencies of a type.

- You can use the `INVALIDATE` option to drop a method that has been redefined, but the redefined versions in the subtypes must still be dropped manually. The subtypes will remain in an invalid state until they are explicitly altered to drop the redefined versions. Until then, an attempt to recompile the subtypes for revalidation will produce the error `Method does not override`.

Unlike `CASCADE`, `INVALIDATE` bypasses all the type and table checks and simply invalidates all schema objects dependent on the type. The objects are revalidated the next time they are accessed. This option is faster than using `CASCADE`, but you must be certain that no problems will be encountered revalidating dependent types and tables. Table data cannot be accessed while a table is invalid; if a table cannot be validated, its data remains inaccessible.

**See Also:** ["If a Type Change Validation Fails"](#) on page 7-15

**Modifying the FINAL or INSTANTIABLE Property**

- Altering a user-defined type from `INSTANTIABLE` to `NOT INSTANTIABLE` is allowed only if the type has no table dependents.
- Altering a user-defined type from `NOT INSTANTIABLE` to `INSTANTIABLE` is allowed anytime. This change does not affect tables.
- Altering a user-defined type from `NOT FINAL` to `FINAL` is allowed only if the target type has no subtypes.
- When you alter a user-defined type from `FINAL` to `NOT FINAL` or vice versa, you must use `CASCADE` to convert data in dependent columns and tables immediately. You may not use the `CASCADE` option `NOT INCLUDING TABLE DATA` to defer converting data.

If you alter a type from `NOT FINAL` to `FINAL`, you must use `CASCADE INCLUDING TABLE DATA`. If you alter a type from `FINAL` to `NOT FINAL`, you may use either `CASCADE INCLUDING TABLE DATA` or `CASCADE CONVERT TO SUBSTITUTABLE`.

When you alter a type from `FINAL` to `NOT FINAL`, the `CASCADE` option you should choose depends on whether you want to be able to insert new subtypes of the type you are altering in existing columns and tables.

By default, altering a type from `FINAL` to `NOT FINAL` enables you to create new substitutable tables and columns of that type, but it does not automatically make existing columns (or object tables) of that type substitutable. In fact, just the opposite happens: existing columns and tables of the type are marked `NOT SUBSTITUTABLE AT ALL LEVELS`. If any embedded attribute of such a column is substitutable, an error is generated. New subtypes of the altered type cannot be inserted in such preexisting columns and tables.

To alter a user-defined type to `NOT FINAL` in such a way as to make existing columns and tables of the type substitutable (assuming that they are not marked `NOT SUBSTITUTABLE`), use the `CASCADE` option `CONVERT TO SUBSTITUTABLE`. For example:

```
CREATE TYPE shape AS OBJECT (  
    name VARCHAR2(30),  
    area NUMBER)  
    FINAL;  
  
/  
ALTER TYPE shape NOT FINAL CASCADE CONVERT TO SUBSTITUTABLE;
```

This `CASCADE` option marks each existing column as `SUBSTITUTABLE AT ALL LEVELS` and causes a new, hidden column to be added for the `Type` of

instances stored in the column. The column can then store subtype instances of the altered type.

## If a Type Change Validation Fails

The `INVALIDATE` option of the `ALTER TYPE` statement lets you alter a type without propagating the type change to dependent objects. In this case, the system does not validate the dependent types and tables to ensure that all the ramifications of the type change are legal. Instead, all dependent schema objects are marked invalid. The objects, including types and tables, are revalidated when next referenced. If a type cannot be revalidated, it remains invalid, and any tables referencing it become inaccessible until the problem is corrected.

A table may fail validation because, for example, adding a new attribute to a type has caused the number of columns in the table to exceed the maximum allowable number of 1000, or because an attribute used as a partitioning or clustering key of a table was dropped from a type.

To force a revalidation of a type, users can issue the `ALTER TYPE COMPILE` statement. To force a revalidation of an invalid table, users can issue the `ALTER TABLE UPGRADE` statement and specify whether the data is to be converted to the latest type version. (Note that, in a table validation triggered by the system when a table is referenced, table data is always updated to the latest type version: you do not have the option to postpone conversion of the data.)

If a table is unable to convert to the latest type version, then `INSERT`, `UPDATE` and `DELETE` statements on the table are not allowed and its data becomes inaccessible. The following DDLs can be executed on the table, but all other statements which reference an invalid table are not allowed until the table is successfully validated:

- `DROP TABLE`
- `TRUNCATE TABLE`

All PL/SQL programs containing variables defined using `%ROWTYPE` of a table or `%TYPE` of a column or attribute from a table are compiled based on the latest type version. If the table fails the revalidation, then compiling any program units that reference that table will also fail.

## ALTER TYPE Statement for Type Evolution

[Table 7-1](#) lists some of the important options in the `ALTER TYPE` statement for altering the attribute or method definition of a type.

**Table 7-1 ALTER TYPE Options for Type Evolution**

Option	Description
INVALIDATE	<p>Invalidates all dependent objects. Using this option bypasses all the type and table checks, to save time.</p> <p>Use this option only if you are certain that problems will not be encountered validating dependent types and tables. Table data cannot be accessed again until it is validated; if it cannot be validated, it remains inaccessible.</p>
CASCADE	<p>Propagates the type change to dependent types and tables. The statement aborts if an error is found in dependent types or tables unless the <code>FORCE</code> option is specified.</p> <p>If <code>CASCADE</code> is specified with no other options, then the <code>INCLUDING TABLE DATA</code> option for <code>CASCADE</code> is implied, and Oracle converts all table data to the latest type version.</p>
INCLUDING TABLE DATA	<p>Converts data stored in all user-defined columns to the most recent version of the column's type.</p> <p>For each new attribute added to the column's type, a new attribute is added to the data and is initialized to <code>NULL</code>. For each attribute dropped from the referenced type, the corresponding attribute data is removed from each row in the table. All tablespaces containing the table's data must be in read write mode; otherwise, the statement will not succeed.</p>

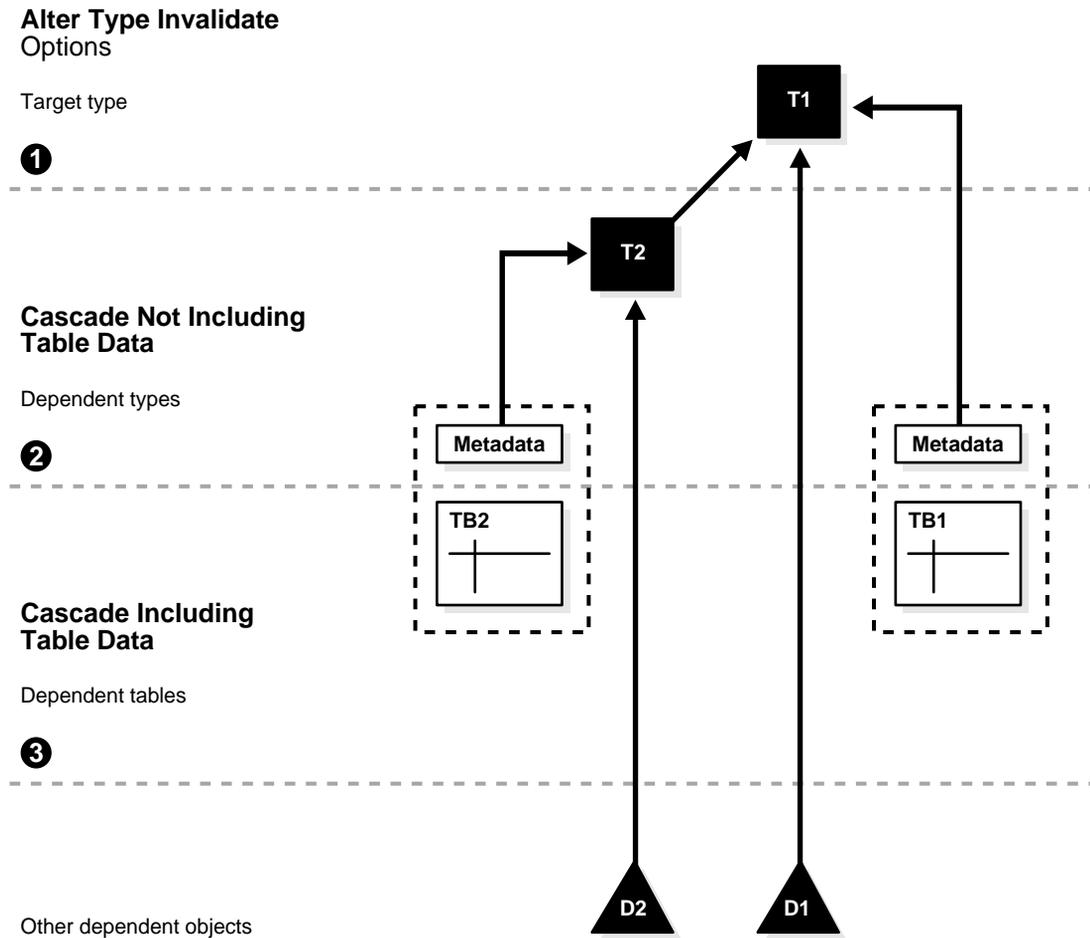
**Table 7–1 (Cont.) ALTER TYPE Options for Type Evolution**

Option	Description
NOT INCLUDING TABLE DATA	<p>Leaves column data as is, associated with the current type version. If an attribute is dropped from a type referenced by a table, then the corresponding column of the dropped attribute is not removed from the table. Only the metadata of the column is marked unused. If the dropped attribute is stored out-of-line (for example, VARRAY, LOB or nested table attribute) then the out-of-line data is not removed. (Unused columns can be removed afterward by using an ALTER TABLE DROP UNUSED COLUMNS statement.)</p> <p>This option is useful when you have many large tables and may run out of rollback segments if you convert them all in one transaction. This option enables you to convert the data of each dependent table later in a separate transaction (using an ALTER TABLE UPGRADE INCLUDING DATA statement).</p> <p>Specifying this option will speed up the table upgrade because the table's data is left in the format of the old type version. However, selecting data from this table will require converting the images stored in the column to the latest type version. This is likely to affect performance during subsequent SELECT statements.</p> <p>Because this option only requires updating the table's metadata all tablespaces are not required to be on-line in read/write mode for the statement to succeed.</p>
FORCE	<p>Forces the system to ignore errors from dependent tables and indexes. Errors are logged in a specified exception table so that they can be queried afterward. This option must be used with caution because dependent tables may become inaccessible if some table errors occur.</p>
CONVERT TO SUBSTITUTABLE	<p>For use when altering a type from FINAL to NOT FINAL: Converts data stored in all user-defined columns to the most recent version of the column's type and then marks these existing columns and object tables of the type SUBSTITUTABLE AT ALL LEVELS so that they can store any new subtypes of the type that are created.</p> <p>If the type is altered to NOT FINAL without specifying this option, existing columns and tables of the type are marked NOT SUBSTITUTABLE AT ALL LEVELS, and new subtypes of the type cannot be stored in them. You will be able to store such subtypes only in columns and tables created after the type was altered.</p>

**See Also:** *Oracle Database SQL Reference* for information about ALTER TYPE options

Figure 7-1 graphically summarizes the options for ALTER TYPE INVALIDATE and their effects. In the figure, T1 is a type and T2 is a dependent type. See the notes following the figure.

Figure 7-1 ALTER TYPE Options



Notes on the figure:

1. **Invalidate:** All objects following line (1) are marked invalid
2. **Cascade Not Including Table Data:** All objects following line (2) are marked invalid. Metadata of all dependent tables are upgraded to the latest type version, but the table data are not converted.
3. **Cascade Including Table Data:** All objects following line (3) are marked invalid. All dependent tables are upgraded to the latest type version, including the table data.

## ALTER TABLE Statement for Type Evolution

You can use `ALTER TABLE` to convert table data to the latest version of referenced types. For an example of converting table data to latest type version, see ["Steps to Change a Type"](#) on page 7-10. See [Table 7-1](#) on page 7-16 for a discussion of the `INCLUDING DATA` option.

**See Also:** *Oracle Database SQL Reference* for information about `ALTER TABLE` options

## The Attribute-Value Constructor

The system-defined attribute value constructor requires you to pass the constructor a value for each attribute of the type. The constructor then sets the attributes of the new object instance to those values. For example:

```
CREATE TYPE shape AS OBJECT (
    name VARCHAR2(30),
    area NUMBER);
/
CREATE TABLE building_blocks OF shape;

-- Attribute value constructor: Sets instance attributes to the specified values
INSERT INTO building_blocks
VALUES (
    NEW shape('my_shape', 4));
```

The keyword `NEW` preceding a call to a constructor is optional but recommended.

## Constructors and Type Evolution

The attribute value constructor function saves you the trouble of defining your own constructors for a type. However, with an attribute-value constructor, you must

supply a value for every attribute declared in the type. Otherwise the constructor call will fail to compile.

This requirement of an attribute-value constructor can create a problem if you evolve the type later on—by adding an attribute, for example. When you change the attributes of a type, the type's attribute-value constructor changes, too. If you add an attribute, the updated attribute-value constructor expects a value for the new attribute as well as the old ones. As a result, all the attribute-value constructor calls in your existing code, where values for only the old number of attributes are supplied, will fail to compile.

**See Also:** ["Type Evolution"](#) on page 7-7

## Advantages of User-Defined Constructors

User-defined constructors avoid the problem with the attribute-value constructor because user-defined constructors do not need to explicitly set a value for every attribute of a type. A user-defined constructor can have any number of arguments, of any type, and these do not need to map directly to type attributes. In your definition of the constructor, you can initialize the attributes to any appropriate values. Any attributes for which you do not supply values are initialized by the system to `NULL`.

If you evolve a type—for example, by adding an attribute—calls to user-defined constructors for the type do not need to be changed. User-defined constructors, like ordinary methods, are not automatically modified when the type evolves, so the call signature of a user-defined constructor remains the same. You may, however, need to change the definition of the constructor if you do not want the new attribute to be initialized to `NULL`.

## Defining and Implementing User-Defined Constructors

You define user-defined constructors in the type body, like an ordinary method function. You introduce the declaration and the definition with the phrase `CONSTRUCTOR FUNCTION`; you must also use the clause `RETURN SELF AS RESULT`.

A constructor for a type must have the same name as the type. [Example 7-1](#) defines two constructor functions for the `shape` type. As the example shows, you can overload user-defined constructors by defining multiple versions with different signatures.

**Example 7–1 Defining and Implementing User-Defined Constructors**

```
CREATE TYPE shape AS OBJECT
(
  name VARCHAR2(30),
  area NUMBER,
  CONSTRUCTOR FUNCTION shape(name VARCHAR2) RETURN SELF AS RESULT,
  CONSTRUCTOR FUNCTION shape(name VARCHAR2, area NUMBER) RETURN
    SELF AS RESULT
) NOT FINAL;
/

CREATE TYPE BODY shape AS
  CONSTRUCTOR FUNCTION shape(name VARCHAR2) RETURN SELF AS RESULT IS
  BEGIN
    SELF.name := name;
    SELF.area := 0;
    RETURN;
  END;
  CONSTRUCTOR FUNCTION shape(name VARCHAR2, area NUMBER) RETURN
  SELF AS RESULT IS
  BEGIN
    SELF.name := name;
    SELF.area := area;
    RETURN;
  END;
END;
/
```

A user-defined constructor has an implicit first parameter `SELF`. Specifying this parameter in the declaration of a user-defined constructor is optional. If you do specify it, its mode must be declared to be `IN OUT`.

The required clause `RETURN SELF AS RESULT` ensures that the most specific type of the instance being returned is the same as the most specific type of the `SELF` argument. In the case of constructors, this is the type for which the constructor is defined.

For example, if the most specific type of the `SELF` argument on a call to the `shape` constructor is `shape`, then this clause ensures that the `shape` constructor returns an instance of `shape` (not an instance of a subtype of `shape`).

When a constructor function is called, the system initializes the attributes of the `SELF` argument to `NULL`. Names of attributes subsequently initialized in the function body may be qualified with `SELF`, as shown in the preceding example, to distinguish them from the names of the arguments of the constructor function, if

these are the same. If the argument names are different, no such qualification is necessary. For example:

```
SELF.name := name;
```

or:

```
name := pl;
```

The function body must include an explicit `return;` as shown. The **return** keyword must not be followed by a `return` expression. The system automatically returns the newly constructed `SELF` instance.

A user-defined constructor may be implemented in PL/SQL, C, or Java.

## Overloading and Hiding Constructors

Like other type methods, user-defined constructors can be overloaded.

User-defined constructors are not inherited, so a user-defined constructor defined in a supertype cannot be hidden in a subtype. However, a user-defined constructor does hide, and thus supersede, the attribute-value constructor for its type if the signature of the user-defined constructor exactly matches the signature of the attribute-value constructor. For the signatures to match, the names and types of the parameters (after the implicit `SELF` parameter) of the user-defined constructor must be the same as the names and types of the type's attributes. The mode of each of the user-defined constructor's parameters (after the implicit `SELF` parameter) must be `IN`.

If an attribute-value constructor is not hidden by a user-defined constructor having the same name and signature, the attribute-value constructor can still be called.

Note that, if you evolve a type—for example, by adding an attribute—the signature of the type's attribute-value constructor changes accordingly. This can cause a formerly hidden attribute-value constructor to become usable again.

## Calling User-Defined Constructors

A user-defined constructor is called like any other function. You can use a user-defined constructor anywhere you can use an ordinary function.

The `SELF` argument is passed in implicitly and may not be passed in explicitly. In other words, usages like the following are not allowed:

```
NEW constructor(instance, argument_list)
```

A user-defined constructor cannot occur in the `DEFAULT` clause of a `CREATE` or `ALTER TABLE` statement, but an attribute-value constructor can. The arguments to the attribute-value constructor must not contain references to PL/SQL functions or to other columns, including the pseudocolumns `LEVEL`, `PRIOR`, and `ROWNUM`, or to date constants that are not fully specified. The same is true for check constraint expressions: an attribute-value constructor can be used as part of check constraint expressions while creating or altering a table, but a user-defined constructor cannot.

Parentheses are required in SQL even for constructor calls that have no arguments. In PL/SQL, parentheses are optional when invoking a zero-argument constructor. They do, however, make it more obvious that the constructor call is a function call. The following PL/SQL example omits parentheses in the constructor call to create a new shape:

```
shape s := NEW my_schema.shape;
```

The `NEW` keyword and the schema name are optional.

### **Example 7-2 Calling User-Defined Constructors with SQL and PL/SQL**

```
CREATE TYPE rectangle UNDER shape
(
    length NUMBER,
    width NUMBER,
    CONSTRUCTOR FUNCTION rectangle(
        name VARCHAR2, length NUMBER, width NUMBER
    ) RETURN SELF AS RESULT);
/

CREATE TYPE BODY rectangle IS
    CONSTRUCTOR FUNCTION rectangle(
        name VARCHAR2, length NUMBER, width NUMBER
    ) RETURN SELF AS RESULT IS
    BEGIN
        SELF.name := name;
        SELF.area := length*width;
        SELF.length := length;
        SELF.width := width;
        RETURN ;
    END;
END;
/

CREATE TABLE shape_table OF shape;
```

```
INSERT INTO shape_table VALUES(shape('shape1', 20));

INSERT INTO shape_table VALUES(rectangle('rectangle', 2, 5));

INSERT INTO shape_table VALUES(rectangle('Quad', 12, 3));
```

The following query selects the rows in the `shape_table`:

```
SELECT VALUE(s) FROM shape_table s;

VALUE(S)(NAME, AREA)
-----
SHAPE('shape1', 20)
RECTANGLE('rectangle', 10, 2, 5)
RECTANGLE('Quad', 36, 12, 3)
```

The following PL/SQL code calls the constructor:

```
s shape := NEW shape('void');
```

## Constructors for SQLJ Object Types

A SQLJ object type is a SQL object type mapped to a Java class. A SQLJ object type has an attribute-value constructor. It can also have user-defined constructors that are mapped to constructors in the referenced Java class.

### **Example 7-3** *Creating a SQLJ Object*

```
CREATE TYPE address AS OBJECT
  EXTERNAL NAME 'university.address'
  LANGUAGE JAVA USING SQLData
(
  street VARCHAR2(100) EXTERNAL NAME 'street',
  city VARCHAR2(50) EXTERNAL NAME 'city',
  state VARCHAR2(50) EXTERNAL NAME 'state',
  zip_code number EXTERNAL NAME 'zipCode',
  CONSTRUCTOR FUNCTION address (full_address VARCHAR)
    EXTERNAL NAME 'address (java.lang.String)');
/
```

A SQLJ type of a serialized representation can have only a user-defined constructor. The internal representation of an object of SQLJ type is opaque to SQL, so an attribute-value constructor is not possible for a SQLJ type.

## Transient and Generic Types

Oracle has three special SQL datatypes that enable you to dynamically encapsulate and access type descriptions, data instances, and sets of data instances of any other SQL type, including object and collection types. You can also use these three special types to create anonymous types, including anonymous collection types.

The three SQL types are implemented as opaque types. In other words, the internal structure of these types is not known to the database; their data can be queried only by implementing functions (typically 3GL routines) for the purpose. Oracle provides both an OCI and a PL/SQL API for implementing such functions.

The three generic SQL types are described in [Table 7-2](#).

**Table 7-2 Generic SQL Types**

Type	Description
<code>SYS.ANYTYPE</code>	<p>A type description type. A <code>SYS.ANYTYPE</code> can contain a type description of any SQL type, named or unnamed, including object types and collection types.</p> <p>An <code>ANYTYPE</code> can contain a type description of a persistent type, but an <code>ANYTYPE</code> itself is transient: in other words, the value in an <code>ANYTYPE</code> itself is not automatically stored in the database. To create a persistent type, use a <code>CREATE TYPE</code> statement from SQL.</p>
<code>SYS.ANYDATA</code>	<p>A self-describing data instance type. A <code>SYS.ANYDATA</code> contains an instance of a given type, with data, plus a description of the type. In this sense, a <code>SYS.ANYDATA</code> is self-describing. An <code>ANYDATA</code> can be persistently stored in the database.</p>
<code>SYS.ANYDATASET</code>	<p>A self-describing data set type. A <code>SYS.ANYDATASET</code> type contains a description of a given type plus a set of data instances of that type. An <code>ANYDATASET</code> can be persistently stored in the database.</p>

Each of these three types can be used with any built-in type native to the database as well as with object types and collection types, both named and unnamed. The types provide a generic way to work dynamically with type descriptions, lone instances, and sets of instances of other types. Using the APIs, you can create a transient `ANYTYPE` description of any kind of type. Similarly, you can create or convert (cast) a data value of any SQL type to an `ANYDATA` and can convert an `ANYDATA` (back) to a SQL type. And similarly again with sets of values and `ANYDATASET`.

The generic types simplify working with stored procedures. You can use the generic types to encapsulate descriptions and data of standard types and pass the encapsulated information into parameters of the generic types. In the body of the procedure, you can detail how to handle the encapsulated data and type descriptions of whatever type.

You can also store encapsulated data of a variety of underlying types in one table column of type ANYDATA or ANYDATASET. For example, you can use ANYDATA with Advanced Queuing to model queues of heterogeneous types of data. You can query the data of the underlying datatypes like any other data.

**Example 7-4** defines and executes a PL/SQL procedure that uses methods built into SYS.ANYDATA to access information about data stored in a SYS.ANYDATA table column.

**Example 7-4 Using SYS.ANYDATA**

```
CREATE OR REPLACE TYPE dogowner AS OBJECT (
    ownerno NUMBER, ownersname VARCHAR2(10) );
/
CREATE OR REPLACE TYPE dog AS OBJECT (
    breed VARCHAR2(10), dogname VARCHAR2(10) );
/
CREATE TABLE mytab ( id NUMBER, data SYS.ANYDATA );
INSERT INTO mytab VALUES ( 1, SYS.ANYDATA.ConvertNumber (5) );
INSERT INTO mytab VALUES ( 2, SYS.ANYDATA.ConvertObject (
    dogowner ( 5555, 'John' ) ) );
commit;

CREATE OR REPLACE procedure P IS
    CURSOR cur IS SELECT id, data FROM mytab;

    v_id mytab.id%TYPE;
    v_data mytab.data%TYPE;
    v_type SYS.ANYTYPE;
    v_typecode PLS_INTEGER;
    v_typername VARCHAR2(60);
    v_dummy PLS_INTEGER;
    v_n NUMBER;
    v_dogowner dogowner;
    non_null_anytype_for_NUMBER exception;
    unknown_typername exception;

BEGIN
    OPEN cur;
```

```

LOOP
  FETCH cur INTO v_id, v_data;
  EXIT WHEN cur%NOTFOUND;
  v_typecode := v_data.GetType ( v_type /* OUT */ );
  CASE v_typecode
    WHEN Dbms_Types.Typecode_NUMBER THEN
      IF v_type IS NOT NULL
        THEN RAISE non_null_anytype_for_NUMBER; END IF;
      v_dummy := v_data.GetNUMBER ( v_n /* OUT */ );
      Dbms_Output.Put_Line (
        To_Char(v_id) || ': NUMBER = ' || To_Char(v_n) );
    WHEN Dbms_Types.Typecode_Object THEN
      v_typename := v_data.GetTypeName();
      IF v_typename NOT IN ( 'OBJTEST1.DOGOWNER' )
        THEN RAISE unknown_typename; END IF;
      v_dummy := v_data.GetObject ( v_dogowner /* OUT */ );
      Dbms_Output.Put_Line (
        To_Char(v_id) || ': user-defined type = ' || v_typename ||
        '(' || v_dogowner.ownerno || ', ' || v_dogowner.ownername || ')' );
    END CASE;
  END LOOP;
  CLOSE cur;

EXCEPTION
  WHEN non_null_anytype_for_NUMBER THEN
    RAISE_Application_Error ( -20000,
      'Paradox: the return AnyType instance FROM GetType ' ||
      'should be NULL for all but user-defined types' );
  WHEN unknown_typename THEN
    RAISE_Application_Error ( -20000,
      'Unknown user-defined type ' || v_typename ||
      ' - program written to handle only OBJTEST1.DOGOWNER' );
END;
/

SELECT t.data.gettypename() FROM mytab t;
SET SERVEROUTPUT ON;
EXEC P;

```

The query and the procedure P in the preceding code sample produce output like the following:

```

T.DATA.GETTYNENAME()
-----
SYS.NUMBER

```

```
OBJTEST1.DOGOWNER
```

```
1: NUMBER = 5
```

```
2: user-defined type = OBJTEST1.DOGOWNER(5555, John )
```

Corresponding to the three generic SQL types are three OCI types that model them. Each has a set of functions for creating and accessing the respective type:

- `OCIType`, corresponding to `SYS.ANYTYPE`
- `OCIAnyData`, corresponding to `SYS.ANYDATA`
- `OCIAnyDataSet`, corresponding to `SYS.ANYDATASET`

**See Also:**

- *Oracle Call Interface Programmer's Guide* for the `OCIType`, `OCIAnyData`, and `OCIAnyDataSet` APIs and details on how to use them.
- *PL/SQL Packages and Types Reference* for information about the interfaces to the `ANYTYPE`, `ANYDATA`, and `ANYDATASET` types and about the `DBMS_TYPES` package, which defines constants for built-in and user-defined types, for use with `ANYTYPE`, `ANYDATA`, and `ANYDATASET`.

## User-Defined Aggregate Functions

Oracle provides a number of pre-defined aggregate functions such as `MAX`, `MIN`, `SUM` for performing operations on a set of records. These pre-defined aggregate functions can be used only with scalar data. However, you can create your own custom implementations of these functions, or define entirely new aggregate functions, to use with complex data—for example, with multimedia data stored using object types, opaque types, and LOBs.

User-defined aggregate functions are used in SQL DML statements just like Oracle's own built-in aggregates. Once such functions are registered with the server, Oracle simply invokes the aggregation routines that you supplied instead of the native ones.

User-defined aggregates can be used with scalar data as well. For example, it may be worthwhile to implement special aggregate functions for working with complex statistical data associated with financial or scientific applications.

User-defined aggregates are a feature of the Extensibility Framework. You implement them using `ODCIAggregate` interface routines.

**See Also:** *Oracle Data Cartridge Developer's Guide* for information on using the `ODCIAggregate` interface routines to implement user-defined aggregate functions

## Partitioning Tables That Contain Oracle Objects

Partitioning addresses the key problem of supporting very large tables and indexes by allowing you to decompose them into smaller and more manageable pieces called partitions. Oracle extends partitioning capabilities by letting you partition tables that contain objects, `REFs`, `varrays`, and nested tables. `Varrays` stored in `LOBs` are `equipartitioned` in a way similar to `LOBs`.

**Example 7-5** partitions the purchase order table along zip codes (`ToZip`), which is an attribute of the `ShipToAddr` embedded column object. For the purposes of this example, the `LineItemList` nested table was made a `varray` to illustrate storage for the partitioned `varray`.

---



---

**Restriction:** Nested tables are allowed in tables that are partitioned; however, the storage table associated with the nested table is not partitioned.

---



---

Assuming that the `LineItemList` is defined as a `varray`:

### **Example 7-5 Partitioning a Table That Contains Objects**

```
CREATE TYPE LineItemList_vartyp as varray(10000) of LineItem_objtyp;
/
CREATE TYPE PurchaseOrder_typ AS OBJECT (
    POno                NUMBER,
    Cust_ref            REF Customer_objtyp,
    OrderDate           DATE,
    ShipDate            DATE,
    OrderForm           BLOB,
    LineItemList        LineItemList_vartyp,
    ShipToAddr          Address_objtyp,

    MAP MEMBER FUNCTION
        ret_value RETURN NUMBER,
    MEMBER FUNCTION
        total_value RETURN NUMBER);
/
CREATE TABLE PurchaseOrders_tab of PurchaseOrder_typ
```

```
LOB (OrderForm) store as (nocache logging)
PARTITION BY RANGE (ShipToAddr.zip)
(PARTITION PurOrderZone1_part
VALUES LESS THAN ('59999')
LOB (OrderForm) store as (
storage (INITIAL 10 MINEXTENTS 10 MAXEXTENTS 100))
VARRAY LineItemList store as LOB (
storage (INITIAL 10 MINEXTENTS 10 MAXEXTENTS 100)),
PARTITION PurOrderZone6_part
VALUES LESS THAN ('79999')
LOB (OrderForm) store as (
storage (INITIAL 10 MINEXTENTS 10 MAXEXTENTS 100))
VARRAY LineItemList store as LOB (
storage (INITIAL 10 MINEXTENTS 10 MAXEXTENTS 100)),
PARTITION PurOrderZone0_part
VALUES LESS THAN ('99999')
LOB (OrderForm) store as (
storage (INITIAL 10 MINEXTENTS 10 MAXEXTENTS 100))
VARRAY LineItemList store as LOB (
storage (INITIAL 10 MINEXTENTS 10 MAXEXTENTS 100)));
```

## How Locators Improve the Performance of Nested Tables

Collection types do not map directly to a native type or structure in languages such as C++ and Java. An application using those languages must access the contents of a collection through Oracle interfaces, such as OCI.

Generally, when the client accesses a nested table explicitly or implicitly (by fetching the containing object), Oracle returns the entire collection value to the client process. For performance reasons, a client may wish to delay or avoid retrieving the entire contents of the collection. Oracle handles this case for you by using a locator instead of the real nested table value. When you really access the contents of the collection, they are automatically transferred to the client.

A nested table locator is like a handle to the collection value. It attempts to preserve the value or copy semantics of the nested table by containing the database snapshot as of its time of retrieval. The snapshot helps the database retrieve the correct instantiation of the nested table value at a later time when the collection elements are fetched using the locator. The locator is scoped to a session and cannot be used across sessions. Because database snapshots are used, it is possible to get a snapshot too old error if there is a high update rate on the nested table. Unlike a LOB locator, the nested table locator is truly a locator and cannot be used to modify the collection instance.

---

# Design Considerations for Oracle Objects

This chapter explains the implementation and performance characteristics of the Oracle object-relational model. Use this information to map a logical data model into an Oracle physical implementation, and when developing applications that use object-oriented features. You should be familiar with the basic concepts behind Oracle objects before you read this chapter.

This chapter covers the following topics:

- [General Storage Considerations for Objects](#)
- [Performance of Object Comparisons](#)
- [Design Considerations for REFs](#)
- [Design Considerations for Collections](#)
- [Design Considerations for Methods](#)
- [Writing Reusable Code Using Invoker Rights](#)
- [Replicating Object Tables and Columns](#)
- [Constraints on Objects](#)
- [Considerations Related to Type Evolution](#)
- [Parallel Queries with Oracle Objects](#)
- [Design Consideration Tips and Techniques](#)

## General Storage Considerations for Objects

This section discusses general storage considerations for various object types.

### Storing Objects as Columns or Rows

You can store objects in columns of relational tables as column objects, or in object tables as row objects. Objects that have meaning outside of the relational database object in which they are contained, or objects that are shared among more than one relational database object, should be made referenceable as row objects. That is, such objects should be stored in an object table instead of in a column of a relational table.

For example, an object of object type `customer` has meaning outside of any particular purchase order, and should be referenceable; therefore, `customer` objects should be stored as row objects in an object table. An object of object type `address`, however, has little meaning outside of a particular purchase order and can be one attribute within a purchase order; therefore, `address` objects should be stored as column objects in columns of relational tables or object tables. So, `address` might be a column object in the `customer` row object.

### Column Object Storage

The storage of a column object is the same as the storage of an equivalent set of scalar columns that collectively make up the object. The only difference is that there is the additional overhead of maintaining the atomic null values of the object and its embedded object attributes. These values are called null indicators because, for every column object, a null indicator specifies whether the column object is null and whether each of its embedded object attributes is null. However, null indicators do not specify whether the scalar attributes of a column object are null. Oracle uses a different method to determine whether scalar attributes are null.

Consider a table that holds the identification number, name, address, and phone numbers of people within an organization. You can create three different object types to hold the name, address, and phone number. First, to create the `name_objtyp` object type, enter the following SQL statement:

```
CREATE TYPE name_objtyp AS OBJECT (  
  first      VARCHAR2(15),  
  middle    VARCHAR2(15),  
  last      VARCHAR2(15));  
/
```

Next, to create the `address_objtyp` object type, enter the following SQL statement:

```
CREATE TYPE address_objtyp AS OBJECT (
  street      VARCHAR2(200),
  city        VARCHAR2(200),
  state       CHAR(2),
  zipcode     VARCHAR2(20));
/
```

Finally, to create the `phone_objtyp` object type, enter the following SQL statement:

```
CREATE TYPE phone_objtyp AS OBJECT (
  location    VARCHAR2(15),
  num         VARCHAR2(14));
/
```

Because each person may have more than one phone number, create a nested table type `phone_ntabtyp` based on the `phone_objtyp` object type:

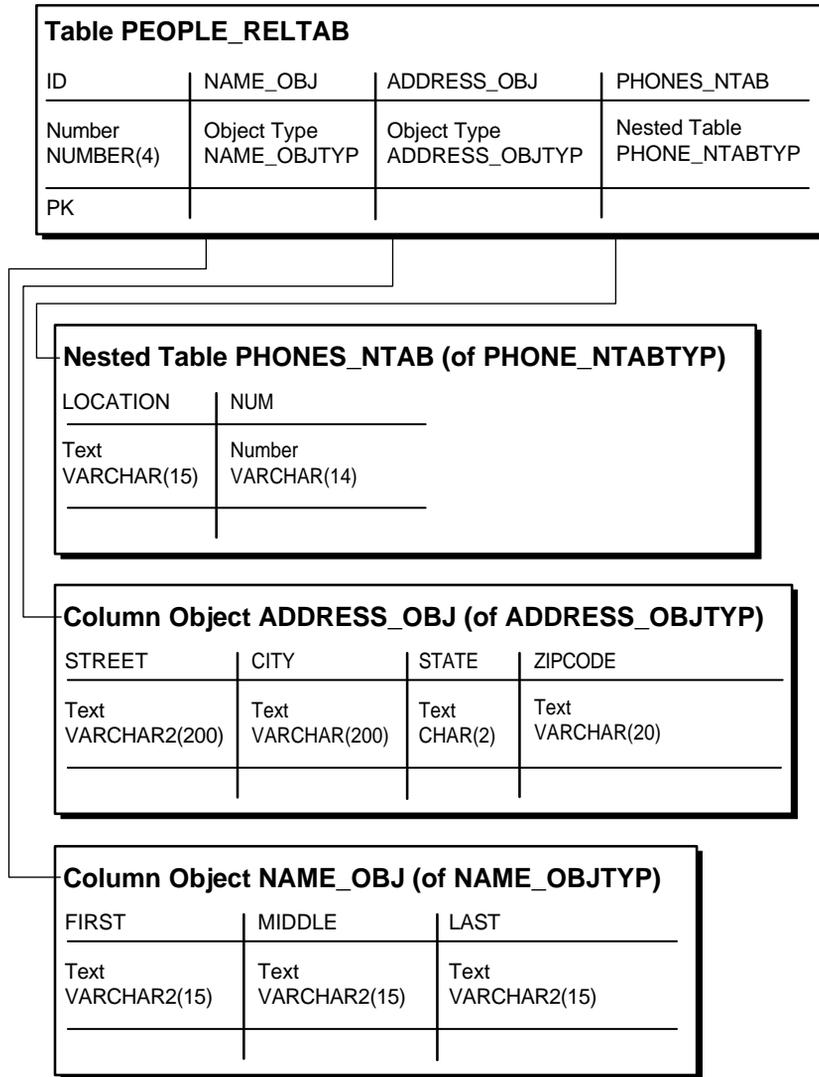
```
CREATE TYPE phone_ntabtyp AS TABLE OF phone_objtyp;
/
```

**See Also:** ["Design Considerations for Nested Tables"](#) on page 8-14 for more information about nested tables.

Once all of these object types are in place, you can create a table to hold the information about the people in the organization with the following SQL statement:

```
CREATE TABLE people_reltab (
  id          NUMBER(4)   CONSTRAINT pk_people_reltab PRIMARY KEY,
  name_obj    name_objtyp,
  address_obj address_objtyp,
  phones_ntab phone_ntabtyp)
NESTED TABLE phones_ntab STORE AS phone_store_ntab;
```

**Figure 8–1 Representation of the people\_reltab Relational Table**



The people\_reltab table has three column objects: name\_obj, address\_obj, and phones\_ntab. The phones\_ntab column object is also a nested table.

---



---

**Note:** The `people_reltab` table and its columns and related types are used in examples throughout this chapter.

---



---

The storage for each object stored in the `people_reltab` table is the same as that of the attributes of the object. For example, the storage required for a `name_obj` object is the same as the storage for the `first`, `middle`, and `last` attributes combined, except for the null indicator overhead.

If the `COMPATIBLE` parameter is set to 8.1.0 or higher, the null indicators for an object and its embedded object attributes occupy one bit each. Thus, an object with  $n$  embedded object attributes (including objects at all levels of nesting) has a storage overhead of  $\text{CEIL}(n/8)$  bytes. In the `people_reltab` table, for example, the overhead of the null information for each row is one byte because it translates to  $\text{CEIL}(3/8)$  or  $\text{CEIL}(.37)$ , which rounds up to one byte. In this case, there are three objects in each row: `name_obj`, `address_obj`, and `phones_ntab`.

If, however, the `COMPATIBLE` parameter is set to a value lower than 8.1.0, such as 8.0.0, the storage is determined by the following calculation:

$$\text{CEIL}(n/8) + 6$$

Here,  $n$  is the total number of all attributes (scalar and object) within the object. Therefore, in the `people_reltab` table, for example, the overhead of the null information for each row is seven bytes because it translates to the following calculation:

$$\text{CEIL}(4/8) + 6 = 7$$

$\text{CEIL}(4/8)$  is  $\text{CEIL}(.5)$ , which rounds up to one byte. In this case, there are three objects in each row and one scalar.

Therefore, the storage overhead and performance of manipulating a column object is similar to that of the equivalent set of scalar columns. The storage for collection attributes are described in the "[Viewing Object Data in Relational Form with Unnesting Queries](#)" section on page 8-11.

**See Also:** *Oracle Database SQL Reference* for more information about `CEIL`.

## Row Object Storage in Object Tables

Row objects are stored in object tables. An object table is a special kind of table that holds objects and provides a relational view of the attributes of those objects. An

object table is logically and physically similar to a relational table whose column types correspond to the top level attributes of the object type stored in the object table. The key difference is that an object table can optionally contain an additional object identifier (OID) column and index.

## Storage Considerations for Object Identifiers (OIDs)

By default, every row object in an object table has an associated logical object identifier (OID) that uniquely identifies it in an object table. Oracle assigns each row object a unique system-generated OID, 16 bytes in length that is automatically indexed for efficient OID-based lookups. The OID column is the equivalent of having an extra 16-byte primary key column. In a distributed and replicated environment, the system-generated unique identifier lets Oracle identify objects unambiguously.

Oracle provides no documentation of or access to the internal structure of object identifiers. This structure can change at any time. The OID column of an object table is a hidden column. Once it is set up, you can ignore it and focus instead on fetching and navigating objects through object references.

An OID allows the corresponding row object to be referred to from other objects or from relational tables. A built-in datatype called a `REF` represents such references. A `REF` encapsulates a reference to a row object of a specified object type.

`REFs` use object identifiers (OIDs) to point to objects. You can use either system-generated OIDs or primary-key based OIDs. The differences between these types of OIDs are outlined in ["Row Object Storage in Object Tables"](#) on page 8-5. If you use system-generated OIDs for an object table, Oracle maintains an index on the column that stores these OIDs. The index requires storage space, and each row object has a system-generated OID, which requires an extra 16 bytes of storage for each row.

### Primary-Key Based OIDs

If a primary key column is available, you can avoid the storage and performance overhead of maintaining the 16-byte OID column and its index. Instead of using the system-generated OIDs, you can use a `CREATE TABLE` statement to specify that the system use the primary key column(s) as the OIDs of the objects in the table. Therefore, you can use existing columns as the OIDs of the objects or use application generated OIDs that are smaller than the 16-byte globally unique OIDs generated by Oracle.

You can avoid these added storage requirements by using the primary key for the object identifiers, instead of system-generated OIDs. You can enforce referential

integrity on columns that store references to these row objects in a way similar to foreign keys in relational tables.

Primary-key based identifiers also make it faster and easier to load data into an object table. By contrast, system-generated object identifiers need to be remapped using some user-specified keys, especially when references to them are also stored.

However, if each primary key value requires more than 16 bytes of storage and you have a large number of `REFs`, using the primary key might require more space than system-generated OIDs because each `REF` is the size of the primary key. In addition, each primary-key based 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 or use system-generated OIDs.

## Performance of Object Comparisons

You can compare objects by invoking the `map` or `order` methods defined on the object type. A `map` method converts objects into scalar values while preserving the ordering of the objects. Mapping objects into scalar values, if it can be done, is preferred because it allows the system to efficiently order objects once they are mapped.

The way objects are mapped has significant performance implications when sorting is required on the objects for `ORDER BY` or `GROUP BY` processing because an object may need to be compared to other objects many times, and it is much more efficient if the objects can be mapped to scalar values first. If the comparison semantics are extremely complex, or if the objects cannot be mapped into scalar values for comparison, you can define an order method that, given two objects, returns the ordering determined by the object implementor. Order methods are not as efficient as `map` methods, so performance may suffer if you use order methods. In any one object type, you can implement either `map` or `order` methods, but not both.

Once again, consider an object type `address` consisting of four character attributes: `street`, `city`, `state`, and `zipcode`. Here, the most efficient comparison method is a `map` method because each object can be converted easily into scalar values. For example, you might define a `map` method that orders all of the objects by `state`.

On the other hand, suppose you want to compare binary objects, such as images. In this case, the comparison semantics may be too complex to use a `map` method; if so, you can use an order method to perform comparisons. For example, you could create an order method that compares images according to brightness or the number of pixels in each image.

If an object type does not have either a map or order method, only equality comparisons are allowed on objects of that type. In this case, Oracle performs the comparison by doing a field-by-field comparison of the corresponding object attributes, in the order they are defined. If the comparison fails at any point, a `FALSE` value is returned. If the comparison matches at every point, a `TRUE` value is returned. However, if an object has a collection of LOB attributes, then Oracle does not compare the object on a field-by-field basis. Such objects must have a map or order method to perform comparisons.

## Design Considerations for REFS

This section discusses considerations when working with REFS.

- [Storage Size of REFS](#)
- [Integrity Constraints for REF Columns](#)
- [Performance and Storage Considerations for Scoped REFS](#)
- [Speeding up Object Access Using the WITH ROWID Option](#)

### Storage Size of REFS

A REF contains the following three logical components:

- OID of the object referenced. A system-generated OID is 16 bytes long. The size of a primary-key based OID depends on the size of the primary key column(s).
- OID of the table or view containing the object referenced, which is 16 bytes long.
- Rowid hint, which is 10 bytes long.

### Integrity Constraints for REF Columns

Referential integrity constraints on REF columns ensure that there is a row object for the REF. Referential integrity constraints on REFS create the same relationship as specifying a primary key/foreign key relationship on relational data. In general, you should use referential integrity constraints wherever possible because they are the only way to ensure that the row object for the REF exists. However, you cannot specify referential integrity constraints on REFS that are in nested tables.

## Performance and Storage Considerations for Scoped REFS

A scoped REF is constrained to contain only references to a specified object table. You can specify a scoped REF when you declare a column type, collection element, or object type attribute to be a REF.

In general, you should use scoped REFS instead of unscoped REFS because scoped REFS are stored more efficiently. Whereas an unscoped REF takes at least 36 bytes to store (more if it uses rowids), a scoped REF is stored as just the OID of its target object and can take less than 16 bytes, depending on whether the referenced OID is system-generated or primary-key based. A system-generated OID requires 16 bytes; a PK-based OID requires enough space to store the primary key value, which may be less than 16 bytes. However, a REF to a PK-based OID, which must be dynamically constructed on being selected, may take more space in memory than a REF to a system-generated OID.

Besides requiring less storage space, scoped REFS often enable the optimizer to optimize queries that dereference a scoped REF into more efficient joins. This optimization is not possible for unscoped REFS because the optimizer cannot determine the containing table(s) for unscoped REFS at query-optimization time.

Unlike referential integrity constraints, scoped REFS do not ensure that the referenced row object exists; they only ensure that the referenced object table exists. Therefore, if you specify a scoped REF to a row object and then delete the row object, the scoped REF becomes a dangling REF because the referenced object no longer exists.

---

---

**Note:** Referential integrity constraints are scoped implicitly.

---

---

Unscoped REFS are useful if the application design requires that the objects referenced be scattered in multiple tables. Because rowid hints are ignored for scoped REFS, you should use unscoped REFS if the performance gain of the rowid hint, as explained in the "[Speeding up Object Access Using the WITH ROWID Option](#)" on page 8-10, outweighs the benefits of the storage saving and query optimization of using scoped REFS.

### Indexing Scoped REFS

You can build indexes on scoped REF columns using the CREATE INDEX command. Then, you can use the index to efficiently evaluate queries that dereference the scoped REFS. Such queries are turned into joins implicitly. For certain types of

queries, Oracle can use an index on the scoped REF column to evaluate the join efficiently.

For example, suppose the object type `address_objtyp` is used to create an object table named `address_objtab`:

```
CREATE TABLE address_objtab OF address_objtyp ;
```

Then, a `people_reltab2` table can be created that has the same definition as the `people_reltab` table discussed in "[Column Object Storage](#)" on page 8-2, except that a REF is used for the address:

```
CREATE TABLE people_reltab2 (  
  id          NUMBER(4)  CONSTRAINT pk_people_reltab2 PRIMARY KEY,  
  name_obj    name_objtyp,  
  address_ref REF address_objtyp SCOPE IS address_objtab,  
  phones_ntab phone_ntabtyp)  
  NESTED TABLE phones_ntab STORE AS phone_store_ntab2 ;
```

Now, an index can be created on the `address_ref` column:

```
CREATE INDEX address_ref_idx ON people_reltab2 (address_ref) ;
```

The following query dereferences the `address_ref`:

```
SELECT id FROM people_reltab2 p  
  WHERE p.address_ref.state = 'CA' ;
```

When this query is executed, the `address_ref_idx` index is used to efficiently evaluate it. Here, `address_ref` is a scoped REF column that stores references to addresses stored in the `address_objtab` object table. Oracle implicitly transforms the preceding query into a query with a join:

```
SELECT p.id FROM people_reltab2 p, address_objtab a  
  WHERE p.address_ref = ref(a) AND a.state = 'CA' ;
```

The Oracle query optimizer might create a plan to perform a nested-loops join with `address_objtab` as the outer table and look up matching addresses using the index on the `address_ref` scoped REF column.

## Speeding up Object Access Using the WITH ROWID Option

If the `WITH ROWID` option is specified for a REF column, Oracle maintains the rowid of the object referenced in the REF. Then, Oracle can find the object referenced directly using the rowid contained in the REF, without the need to fetch the rowid from the OID index. Therefore, you use the `WITH ROWID` option to specify a rowid

hint. Maintaining the rowid requires more storage space because the rowid adds 10 bytes to the storage requirements of the REF.

Bypassing the OID index search improves the performance of REF traversal (navigational access) in applications. The actual performance gain may vary from application to application depending on the following factors:

- How large the OID indexes are.
- Whether the OID indexes are cached in the buffer cache.
- How many REF traversals an application does.

The WITH ROWID option is only a hint because, when you use this option, Oracle checks the OID of the row object with the OID in the REF. If the two OIDs do not match, Oracle uses the OID index instead. The rowid hint is not supported for scoped REFS, for REFS with referential integrity constraints, or for primary key-based REFS.

## Design Considerations for Collections

This section discusses considerations when working with collections.

- [Viewing Object Data in Relational Form with Unnesting Queries](#)
- [Storage Considerations for Varrays](#)
- [Performance of Varrays Versus Nested Tables](#)
- [Design Considerations for Nested Tables](#)
- [Design Considerations for Multilevel Collections](#)

### Viewing Object Data in Relational Form with Unnesting Queries

An unnesting query on a collection allows the data to be viewed in a flat (relational) form. You can execute unnesting queries on single-level and multilevel collections of either nested tables or varrays. This section contains examples of unnesting queries.

Nested tables can be unnested for queries using the TABLE syntax, as in the following example:

```
SELECT p.name_obj, n.num
       FROM people_reltab p, TABLE(p.phones_ntab) n ;
```

Here, `phones_ntab` specifies the attributes of the `phones_ntab` nested table. To retrieve even parent rows that have no child rows (no phone numbers, in this case), use the outer join syntax, with the `+`. For example:

```
SELECT p.name_obj, n.num
       FROM people_reltab p, TABLE(p.phones_ntab) (+) n ;
```

If the `SELECT` list of a query does not refer to any columns from the parent table other than the nested table column, the query is optimized to execute only against the nested table's storage table.

The unnesting query syntax is the same for varrays as for nested tables. For instance, suppose the `phones_ntab` nested table is instead a varray named `phones_var`. The following example shows how to use the `TABLE` syntax to query the varray:

```
SELECT p.name_obj, v.num
       FROM people_reltab p, TABLE(p.phones_var) v ;
```

## Using Procedures and Functions in Unnesting Queries

You can create procedures and functions that you can then execute to perform unnesting queries. For example, you can create a function called `home_phones()` that returns only the phone numbers where `location` is `home`. To create the `home_phones()` function, you enter code like the following:

```
CREATE OR REPLACE FUNCTION home_phones(allphones IN phone_ntabtyp)
    RETURN phone_ntabtyp IS
    homephones phone_ntabtyp := phone_ntabtyp();
    indx1      number;
    indx2      number := 0;
BEGIN
    FOR indx1 IN 1..allphones.count LOOP
        IF
            allphones(indx1).location = 'home'
        THEN
            homephones.extend;    -- extend the local collection
            indx2 := indx2 + 1;
            homephones(indx2) := allphones(indx1);
        END IF;
    END LOOP;

    RETURN homephones;
END;
```

Now, to query for a list of people and their home phone numbers, enter the following:

```
SELECT p.name_obj, n.num
       FROM people_reltab p, TABLE(
           CAST(home_phones(p.phones_ntab) AS phone_ntabtyp)) n ;
```

To query for a list of people and their home phone numbers, including those people who do not have a home phone number listed, enter the following:

```
SELECT p.name_obj, n.num
       FROM people_reltab p,
           TABLE(home_phones(p.phones_ntab) AS phone_ntabtyp)(+) n ;
```

**See Also:** *Oracle Database SQL Reference* for more information about using the TABLE syntax.

## Storage Considerations for Varrays

The size of a stored varray depends only on the current count of the number of elements in the varray and not on the maximum number of elements that it can hold. Because the storage of varrays incurs some overhead, such as null information, the size of the varray stored may be slightly greater than the size of the elements multiplied by the count.

Varrays are stored in columns either as raw values or LOBs. Oracle decides how to store the varray when the varray is defined, based on the maximum possible size of the varray computed using the LIMIT of the declared varray. If the size exceeds approximately 4000 bytes, then the varray is stored in LOBs. Otherwise, the varray is stored in the column itself as a raw value. In addition, Oracle supports inline LOBs which means that elements that fit in the first 4000 bytes of a large varray, with some bytes reserved for the LOB locator, are stored in the column of the row.

### Propagating VARRAY Size Change

When changing the size of a VARRAY type, a new type version is generated for the dependent types. It is important to be aware of this when a VARRAY column is not explicitly stored as a LOB and its maximum size is originally smaller than 4000 bytes. If the size is larger than or equal to 4000 bytes after the increase, the VARRAY column has to be stored as a LOB. This requires an extra operation to upgrade the metadata of the VARRAY column in order to set up the necessary LOB metadata information including the LOB segment and LOB index.

The `CASCADE` option in the `ALTER TYPE` statement propagates the `VARRAY` size change to its dependent types and tables. A new version is generated for each valid dependent type and dependent tables metadata are updated accordingly based on the different case scenarios described previously. If the `VARRAY` column is in a cluster table, an `ALTER TYPE` statement with the `CASCADE` option fails because a cluster table does not support a LOB.

The `CASCADE` option in the `ALTER TYPE` statement also provides the `[NOT] INCLUDING TABLE DATA` option. The `NOT INCLUDING TABLE DATA` option only updates the metadata of the table, but does not convert the data image. In order to convert the `VARRAY` image to the latest version format, you can either specify `INCLUDING TABLE DATA` explicitly in `ALTER TYPE CASCADE` statement or issue `ALTER TABLE UPGRADE` statement.

## Performance of Varrays Versus Nested Tables

If the entire collection is manipulated as a single unit in the application, varrays perform much better than nested tables. The varray is stored packed and requires no joins to retrieve the data, unlike nested tables.

### Varray Querying

The unnesting syntax can be used to access varray columns similar to the way it is used to access nested tables.

**See Also:** ["Viewing Object Data in Relational Form with Unnesting Queries"](#) on page 8-11 for more information.

### Varray Updates

Piece-wise updates of a varray value are not supported. Thus, when a varray is updated, the entire old collection is replaced by the new collection.

## Design Considerations for Nested Tables

The following sections contain design considerations for using nested tables.

### Nested Table Storage

Oracle stores the rows of a nested table in a separate storage table. A system generated `NESTED_TABLE_ID`, which is 16 bytes in length, correlates the parent row with the rows in its corresponding storage table.

Figure 8–2 shows how the storage table works. The storage table contains each value for each nested table in a nested table column. Each value occupies one row in the storage table. The storage table uses the `NESTED_TABLE_ID` to track the nested table for each value. So, in Figure 8–2, all of the values that belong to nested table A are identified, all of the values that belong to nested table B are identified, and so on.

**Figure 8–2 Nested Table Storage**

DATA1	DATA2	DATA3	DATA4	NT_DATA
...	...	...	...	A
...	...	...	...	B
...	...	...	...	C
...	...	...	...	D
...	...	...	...	E

Storage Table

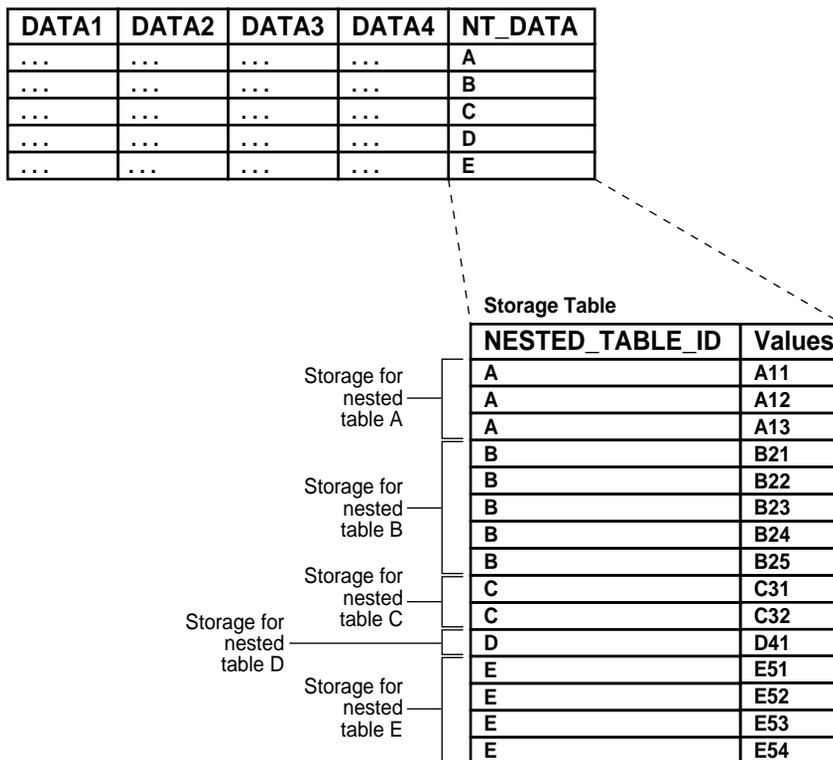
NESTED_TABLE_ID	Values
B	B21
B	B22
C	C33
A	A11
E	E51
B	B25
E	E52
A	A12
E	E54
B	B23
C	C32
A	A13
D	D41
B	B24
E	E53

### Nested Table in an Index-Organized Table (IOT)

If a nested table has a primary key, you can organize the nested table as an index-organized table (IOT). If the `NESTED_TABLE_ID` column is a prefix of the primary key for a given parent row, Oracle physically clusters its child rows together. So, when a parent row is accessed, all its child rows can be efficiently retrieved. When only parent rows are accessed, efficiency is maintained because the child rows are not inter-mixed with the parent rows.

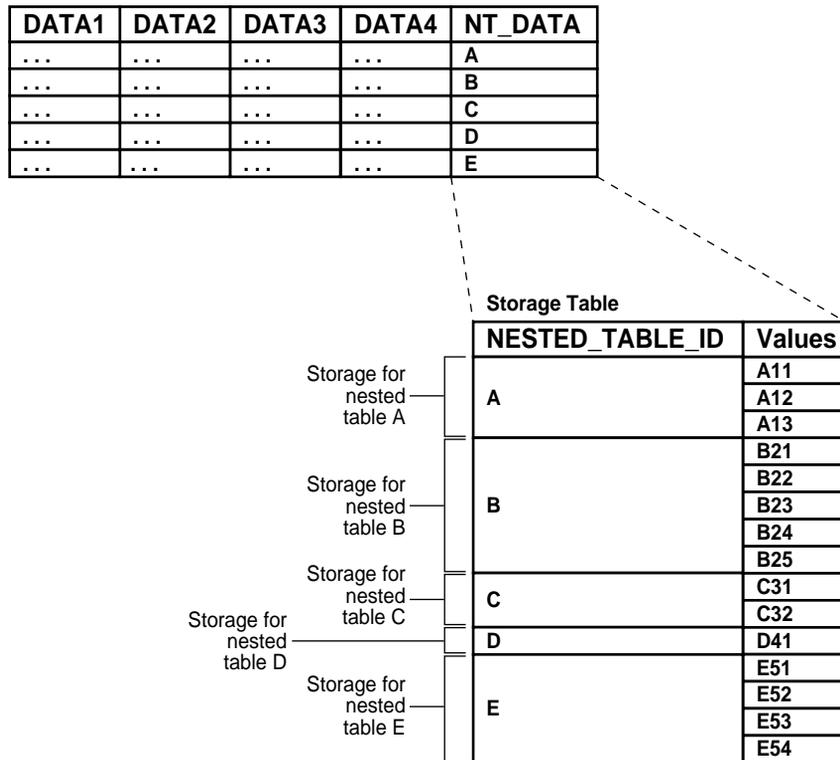
Figure 8–3 shows how the storage table works when the nested table is in an IOT. The storage table groups by `NESTED_TABLE_ID` the values for each nested table in a nested table column. In Figure 8–3, for each nested table in the `NT_DATA` column of the parent table, the data is grouped in the storage table: all of the values in nested table A are grouped together, all of the values in nested table B are grouped together, and so on.

**Figure 8–3 Nested Table in IOT Storage**



In addition, the `COMPRESS` clause enables prefix compression on the IOT rows. It factors out the key of the parent in every child row. That is, the parent key is not repeated in every child row, thus providing significant storage savings.

In other words, if you specify nested table compression using the `COMPRESS` clause, the amount of space required for the storage table is reduced because the `NESTED_TABLE_ID` is not repeated for each value in a group. Instead, the `NESTED_TABLE_ID` is stored only once for each group, as illustrated in Figure 8–4.

**Figure 8–4 Nested Table in IOT Storage with Compression**

In general, Oracle recommends that nested tables be stored in an IOT with the `NESTED_TABLE_ID` column as a prefix of the primary key. Further, prefix compression should be enabled on the IOT. However, if you usually do not retrieve the nested table as a unit and you do not want to cluster the child rows, do not store the nested table in an IOT and do not specify compression.

### Nested Table Indexes

For nested tables stored in heap tables (as opposed to IOTs), you should create an index on the `NESTED_TABLE_ID` column of the storage table. The index on the corresponding ID column of the parent table is created by Oracle automatically when the table is created. Creating an index on the `NESTED_TABLE_ID` column enables Oracle to access the child rows of the nested table more efficiently, because Oracle must perform a join between the parent table and the nested table using the `NESTED_TABLE_ID` column.

## Nested Table Locators

For large child sets, the parent row and a locator to the child set can be returned so that the child rows can be accessed on demand; the child sets also can be filtered. Using nested table locators enables you to avoid unnecessarily transporting child rows for every parent.

You can perform either one of the following actions to access the child rows using the nested table locator:

- Call the OCI collection functions. This action occurs implicitly when you access the elements of the collection in the client-side code, such as `OCI*Coll*` functions. The entire collection is retrieved implicitly on the first access.

**See Also:** *Oracle Call Interface Programmer's Guide* for more information about OCI collection functions.

- Use SQL to retrieve the rows corresponding to the nested table.

In a multilevel collection, you can use a locator with a specified collection at any level of nesting. Following are described two ways in which to specify that a collection is to be retrieved as a locator.

## At Table Creation Time

When the collection type is being used as a column type and the `NESTED TABLE` storage clause is used, you can use the `RETURN AS LOCATOR` clause to specify that a particular collection is to be retrieved as a locator.

For instance, suppose that `inner_table` is a collection type consisting of three levels of nested tables. In the following example, the `RETURN AS LOCATOR` clause specifies that the third level of nested tables is always to be retrieved as a locator.

```
CREATE TYPE inner_table AS TABLE OF NUMBER;
/
CREATE TYPE middle_table AS TABLE OF inner_table;
/
CREATE TYPE outer_table AS TABLE OF middle_table;
/
CREATE TABLE tabl (
  col1 NUMBER,
  col2 outer_table)
NESTED TABLE col2 STORE AS col2_ntab
(NESTED TABLE COLUMN_VALUE STORE AS cval1_ntab
(NESTED TABLE COLUMN_VALUE STORE AS cval2_ntab RETURN AS LOCATOR) );
```

### As a HINT During Retrieval

A query can retrieve a collection as a locator by means of the hint `NESTED_TABLE_GET_REFS`. Here is an example of retrieving the column `col2` from the table `tab1` as a locator:

```
SELECT /*+ NESTED_TABLE_GET_REFS */ col2
      FROM tab1
     WHERE col1 = 2;
```

Unlike with the `RETURN AS LOCATOR` clause, however, you cannot specify a particular inner collection to return as a locator when using the hint.

### Optimizing Set Membership Queries

Set membership queries are useful when you want to search for a specific item in a nested table. For example, the following query tests the membership in a child-set; specifically, whether the location `home` is in the nested table `phones_ntab`, which is in the parent table `people_reltab`:

```
SELECT * FROM people_reltab p
       WHERE 'home' IN (SELECT location FROM TABLE(p.phones_ntab)) ;
```

Oracle can execute a query that tests the membership in a child-set more efficiently by transforming it internally into a semijoin. However, this optimization only happens if the `ALWAYS_SEMI_JOIN` initialization parameter is set. If you want to perform semijoins, the valid values for this parameter are `MERGE` and `HASH`; these parameter values indicate which join method to use.

---



---

**Note:** In the preceding example, `home` and `location` are child set elements. If the child set elements are object types, they must have a `map` or `order` method to perform a set membership query.

---



---

## Design Considerations for Multilevel Collections

[Chapter 2](#) describes how to nest collection types to create a true multilevel collection, such as a nested table of nested tables, a nested table of varrays, a varray of nested tables, a varray of nested tables, or a varray or nested table of an object type that has an attribute of a collection type.

You can also nest collections indirectly using `REFS`. For example, you can create a nested table of an object type that has an attribute that references an object that has a nested table or varray attribute. If you do not actually need to access all elements

of a multilevel collection, then nesting a collection with `REFs` may provide better performance because only the `REFs` need to be loaded, not the elements themselves.

True multilevel collections (specifically multilevel nested tables) perform better for queries that access individual elements of the collection. Using nested table locators can improve the performance of programmatic access if you do not need to access all elements.

For an example of a collection that uses `REFs` to nest another collection, suppose you want to create a new object type called `person_objtyp` using the object types described in "[Column Object Storage](#)" on page 8-2, which are `name_objtyp`, `address_objtyp`, and `phone_ntabtyp`. Remember that the `phone_ntabtyp` object type is a nested table because each person may have more than one phone number.

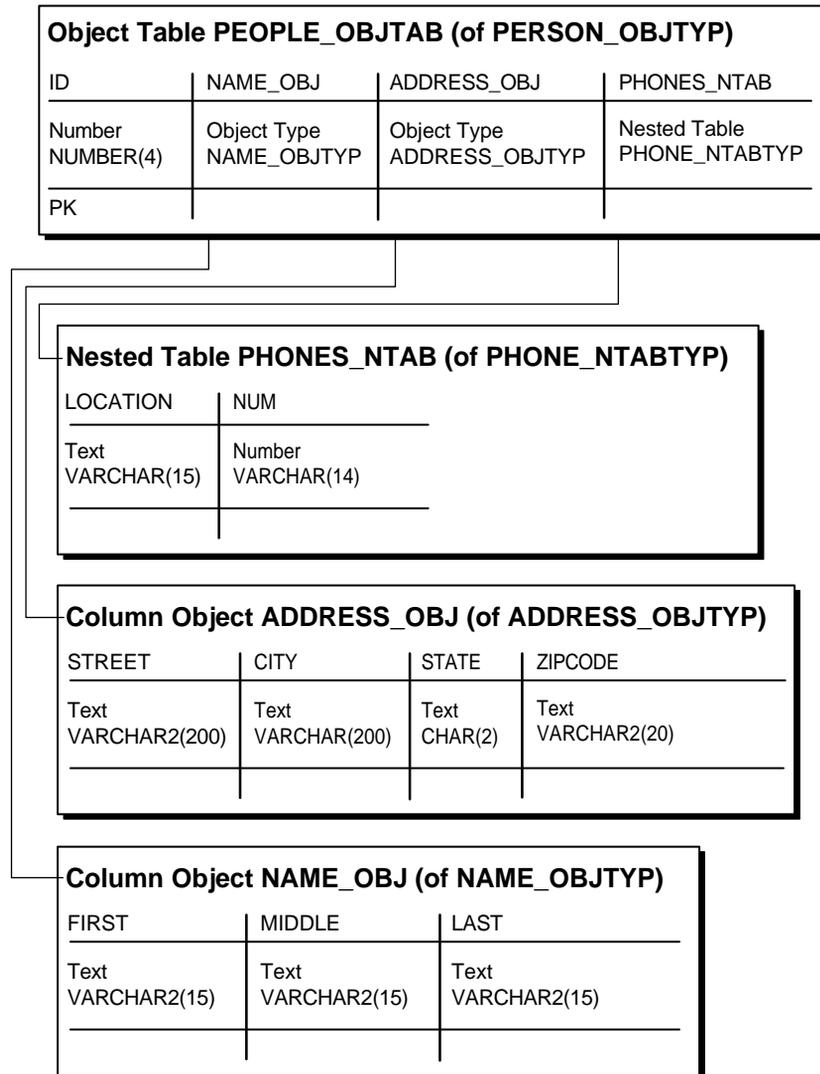
To create the `person_objtyp` object type, issue the following SQL statement:

```
CREATE TYPE person_objtyp AS OBJECT (  
    id          NUMBER(4),  
    name_obj    name_objtyp,  
    address_obj address_objtyp,  
    phones_ntab phone_ntabtyp);
```

To create an object table called `people_objtab` of `person_objtyp` object type, issue the following SQL statement:

```
CREATE TABLE people_objtab OF person_objtyp (id PRIMARY KEY)  
    NESTED TABLE phones_ntab STORE AS phones_store_ntab ;
```

The `people_objtab` table has the same attributes as the `people_reltab` table discussed in "[Column Object Storage](#)" on page 8-2. The difference is that the `people_objtab` is an object table with row objects, while the `people_reltab` table is a relational table with three column objects.

**Figure 8–5 Object-Relational Representation of the people\_objtab Object Table**

Now you can reference the row objects in the `people_objtab` object table from other tables. For example, suppose you want to create a `projects_objtab` table that contains:

- A project identification number for each project.

- The title of each project.
- The project lead for each project.
- A description of each project.
- Nested table collection of the team of people assigned to each project.

You can use REFS to the `people_objtab` for the project leads, and you can use a nested table collection of REFS for the team. To begin, create a nested table object type called `personref_ntabtyp` based on the `person_objtyp` object type:

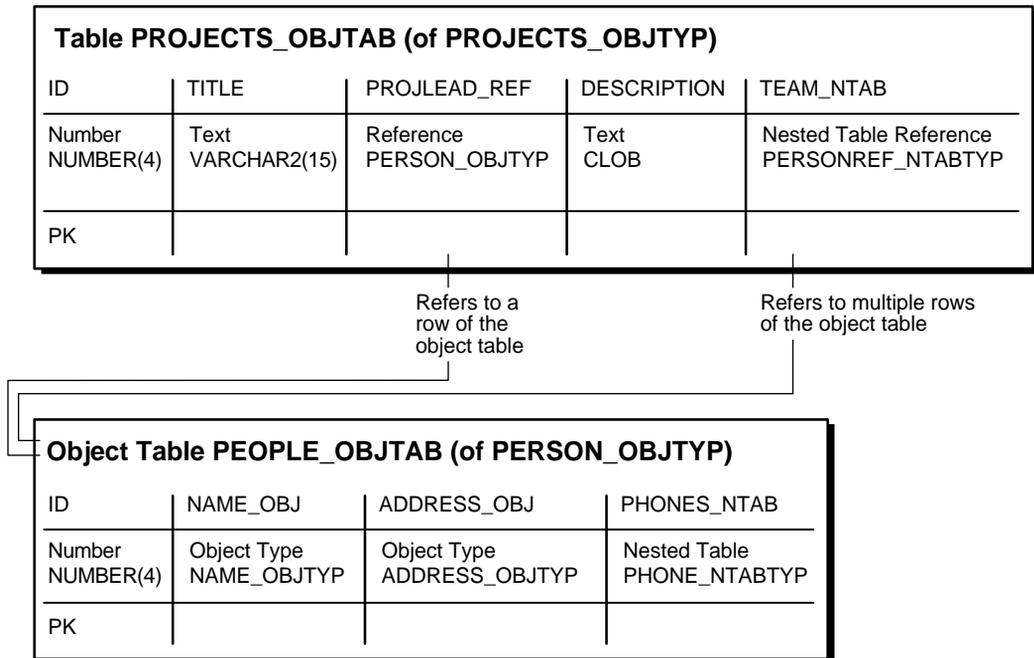
```
CREATE TYPE personref_ntabtyp AS TABLE OF REF person_objtyp;  
/
```

Now you are ready to create the object table `projects_objtab`. First, create the object type `projects_objtyp` by issuing the following SQL statement:

```
CREATE TYPE projects_objtyp AS OBJECT (  
    id            NUMBER(4),  
    title         VARCHAR2(15),  
    projlead_ref REF person_objtyp,  
    description   CLOB,  
    team_ntab    personref_ntabtyp);  
/
```

Next, create the object table `projects_objtab` based on the `projects_objtyp`:

```
CREATE TABLE projects_objtab OF projects_objtyp (id PRIMARY KEY)  
    NESTED TABLE team_ntab STORE AS team_store_ntab ;
```

**Figure 8–6 Object-Relational Representation of the projects\_objtab Object Table**

Once the `people_objtab` object table and the `projects_objtab` object table are in place, you indirectly have a nested collection. That is, the `projects_objtab` table contains a nested table collection of `REFs` that point to the people in the `people_objtab` table, and the people in the `people_objtab` table have a nested table collection of phone numbers.

You can insert values into the `people_objtab` table in the following way:

```
INSERT INTO people_objtab VALUES (
    0001,
    name_objtyp('JOHN', 'JACOB', 'SCHMIDT'),
    address_objtyp('1252 Maple Road', 'Fairfax', 'VA', '22033'),
    phone_ntabtyp(
        phone_objtyp('home', '650.339.9922'),
        phone_objtyp('work', '510.563.8792')));

INSERT INTO people_objtab VALUES (
    0002,
    name_objtyp('MARY', 'ELLEN', 'MILLER'),
    address_objtyp('33 Spruce Street', 'McKees Rocks', 'PA', '15136'),
```

```
phone_ntabtyp(
  phone_objtyp('home', '415.642.6722'),
  phone_objtyp('work', '650.891.7766')) ;

INSERT INTO people_objtab VALUES (
  0003,
  name_objtyp('SARAH', 'MARIE', 'SINGER'),
  address_objtyp('525 Pine Avenue', 'San Mateo', 'CA', '94403'),
  phone_ntabtyp(
    phone_objtyp('home', '510.804.4378'),
    phone_objtyp('work', '650.345.9232'),
    phone_objtyp('cell', '650.854.9233')) ;
```

Then, you can insert into the `projects_objtab` relational table by selecting from the `people_objtab` object table using a `REF` operator, as in the following examples:

```
INSERT INTO projects_objtab VALUES (
  1101,
  'Demo Product',
  (SELECT REF(p) FROM people_objtab p WHERE id = 0001),
  'Demo the product, show all the great features.',
  personref_ntabtyp(
    (SELECT REF(p) FROM people_objtab p WHERE id = 0001),
    (SELECT REF(p) FROM people_objtab p WHERE id = 0002),
    (SELECT REF(p) FROM people_objtab p WHERE id = 0003))) ;

INSERT INTO projects_objtab VALUES (
  1102,
  'Create PRODDB',
  (SELECT REF(p) FROM people_objtab p WHERE id = 0002),
  'Create a database of our products.',
  personref_ntabtyp(
    (SELECT REF(p) FROM people_objtab p WHERE id = 0002),
    (SELECT REF(p) FROM people_objtab p WHERE id = 0003))) ;
```

---

---

**Note:** This example uses nested tables to store `REFs`, but you also can store `REFs` in `varrays`. That is, you can have a `varray` of `REFs`.

---

---

## Design Considerations for Methods

This section discusses considerations when working with methods.

- [Choosing a Language for Method Functions](#)
- [Static Methods](#)
- [Function-Based Indexes on the Return Values of Type Methods](#)

## Choosing a Language for Method Functions

Method functions can be implemented in any of the languages supported by Oracle, such as PL/SQL, Java, or C. Consider the following factors when you choose the language for a particular application:

- Ease of use
- SQL calls
- Speed of execution
- Same/different address space

In general, if the application performs intense computations, C is preferable, but if the application performs a relatively large number of database calls, PL/SQL or Java is preferable.

A method implemented in C executes in a separate process from the server using external procedures. In contrast, a method implemented in Java or PL/SQL executes in the same process as the server.

### Example: Implementing a Method

The example described in this section involves an object type whose methods are implemented in different languages. In the example, the object type `ImageType` has an `ID` attribute, which is a `NUMBER` that uniquely identifies it, and an `IMG` attribute, which is a `BLOB` that stores the raw image. The object type `ImageType` has the following methods:

- The method `get_name` fetches the name of the image by looking it up in the database. This method is implemented in PL/SQL.
- The method `rotate` rotates the image. This method is implemented in C.
- The method `clear` returns a new image of the specified color. This method is implemented in Java.

For implementing a method in C, a `LIBRARY` object must be defined to point to the library that contains the external C routines. For implementing a method implemented in Java, this example assumes that the Java class with the method has been compiled and uploaded into Oracle.

Here is the object type specification and its methods:

```
CREATE LIBRARY myCfuncs TRUSTED AS STATIC
/

CREATE TYPE ImageType AS OBJECT (
    id NUMBER,
    img BLOB,
    MEMBER FUNCTION get_name return VARCHAR2,
    MEMBER FUNCTION rotate return BLOB,
    STATIC FUNCTION clear(color NUMBER) return BLOB);
/

CREATE TYPE BODY ImageType AS
    MEMBER FUNCTION get_name RETURN VARCHAR2
    IS
        imgname VARCHAR2(100);
        sqlstmt VARCHAR2(200);
    BEGIN
        sqlstmt := 'SELECT name INTO imgname FROM imgtab WHERE imgid = id';
        EXECUTE IMMEDIATE sqlstmt;
        RETURN imgname;
    END;

    MEMBER FUNCTION rotate RETURN BLOB
    AS LANGUAGE C
    NAME "Crotate"
    LIBRARY myCfuncs;

    STATIC FUNCTION clear(color NUMBER) RETURN BLOB
    AS LANGUAGE JAVA
    NAME 'myJavaClass.clear(oracle.sql.NUMBER) return oracle.sql.BLOB';

END;
/
```

---

---

**Restriction:** Type methods can be mapped only to static Java methods.

---

---

**See Also:**

- *Oracle Database Java Developer's Guide* for more information
- [Chapter 4, "Object Support in Oracle Programming Environments"](#) for more information about choosing a language

## Static Methods

Static methods differ from member methods in that the `SELF` value is not passed in as the first parameter. Methods in which the value of `SELF` is not relevant should be implemented as static methods. Static methods can be used for user-defined constructors.

The following example is a constructor-like method that constructs an instance of the type based on the explicit input parameters and inserts the instance into the specified table:

```
CREATE TYPE atype AS OBJECT(
  a1 NUMBER,
  STATIC PROCEDURE newa (
    p1          NUMBER,
    tabname     VARCHAR2,
    schname     VARCHAR2));
/
CREATE TYPE BODY atype AS
  STATIC PROCEDURE newa (p1 NUMBER, tabname VARCHAR2, schname VARCHAR2)
  IS
    sqlstmt VARCHAR2(100);
  BEGIN
    sqlstmt := 'INSERT INTO '||schname||'.'||tabname||' VALUES (atype(:1))';
    EXECUTE IMMEDIATE sqlstmt USING p1;
  END;
END;
/

CREATE TABLE atab OF atype;

BEGIN
  atype.newa(1, 'atab', 'OBJTEST1');
END;
/
```

## Using SELF IN OUT NOCOPY with Member Procedures

In member procedures, if `SELF` is not declared, its parameter mode defaults to `IN OUT`. However, the default behavior does not include the `NOCOPY` compiler hint.

Because the value of the `IN OUT` actual parameter is copied into the corresponding formal parameter, the copying slows down execution when the parameters hold large data structures such as instances of large object types.

For performance reasons, you may want to include `SELF IN OUT NOCOPY` when passing a large object type as a parameter. For example:

```
MEMBER PROCEDURE my_proc (SELF IN OUT NOCOPY my_LOB)
```

### See Also:

- *PL/SQL User's Guide and Reference* for information on performance issues and restrictions on the use of `NOCOPY`
- *Oracle Database SQL Reference* for information about using `NOCOPY` in the `CREATE PROCEDURE` statement
- ["Member Methods"](#) on page 2-9

## Function-Based Indexes on the Return Values of Type Methods

A function-based index is an index based on the return values of an expression or function. The function may be a method function of an object type.

A function-based index built on a method function precomputes the return value of the function for each object instance in the column or table being indexed and stores those values in the index. There they can be referenced without having to evaluate the function again.

Function-based indexes are useful for improving the performance of queries that have a function in the `WHERE` clause. For example, the following code contains a query of an object table `emps`:

```
CREATE TYPE emp_t AS OBJECT(  
  name  VARCHAR2(36),  
  salary NUMBER,  
  MEMBER FUNCTION bonus RETURN NUMBER DETERMINISTIC);  
/  
CREATE TYPE BODY emp_t IS  
  MEMBER FUNCTION bonus RETURN NUMBER DETERMINISTIC IS  
  BEGIN
```

```

    RETURN self.salary * .1;
END;
END;
/
CREATE TABLE emps OF emp_t ;

SELECT e.name
       FROM emps e
       WHERE e.bonus() > 2000;

```

To evaluate this query, Oracle must evaluate `bonus()` for each row object in the table. If there is a function-based index on the return values of `bonus()`, then this work has already been done, and Oracle can simply look up the results in the index. This enables Oracle to return a result from the query more quickly.

Return values of a function can be usefully indexed only if those values are constant, that is, only if the function always returns the same value for each object instance. For this reason, to use a user-written function in a function-based index, the function must have been declared with the `DETERMINISTIC` keyword, as in the preceding example. This keyword promises that the function always returns the same value for each object instance's set of input argument values.

The following example creates a function-based index on the method `bonus()` in the table `emps`:

```
CREATE INDEX emps_bonus_idx ON emps x (x.bonus());
```

**See Also:** *Oracle Database Concepts* and *Oracle Database SQL Reference* for detailed information about function-based indexes

## Writing Reusable Code Using Invoker Rights

To create generic object types that can be used in any schema, you must define the type to use invoker rights, through the `AUTHID CURRENT_USER` option of `CREATE` or `REPLACE TYPE`. In general, use invoker rights when both of the following conditions are true:

- There are type methods that access and manipulate data.
- Users who did not define these type methods must use them.

For example, you can grant user `OBJTEST2` execute privileges on type `atype` created by `OBJTEST1` in "[Static Methods](#)" on page 8-27, and then create table `atab` based on the type:

```
GRANT EXECUTE ON atype TO OBJTEST2;
CONNECT OBJTEST2/OBJTEST2;
CREATE TABLE atab OF OBJTEST1.atype ;
```

Now, suppose user OBJTEST2 tries to use atype in the following statement:

```
BEGIN
  OBJTEST1.atype.newa(1, 'atab', 'OBJTEST2'); -- raises an error
END;
/
```

**This statement raises an error because the definer of the type (OBJTEST1) does not have the privileges required to perform the insert in the newa procedure. You can avoid this error by defining atype using invoker rights. Here, you first drop the atab table in both schemas and re-create atype using invoker rights:**

```
DROP TABLE atab;
CONNECT OBJTEST1/OBJTEST1;
DROP TABLE atab; DROP TYPE atype; commit;

CREATE TYPE atype AUTHID CURRENT_USER AS OBJECT(
  a1 NUMBER,
  STATIC PROCEDURE newa(p1 NUMBER, tabname VARCHAR2, schname VARCHAR2));
/
CREATE TYPE BODY atype AS
  STATIC PROCEDURE newa(p1 NUMBER, tabname VARCHAR2, schname VARCHAR2)
  IS
    sqlstmt VARCHAR2(100);
  BEGIN
    sqlstmt := 'INSERT INTO '||schname||'.'||tabname|| '
              VALUES (OBJTEST1.atype(:1))';
    EXECUTE IMMEDIATE sqlstmt USING p1;
  END;
END;
/
```

**Now, if user OBJTEST2 tries to use atype again, the statement executes successfully:**

```
GRANT EXECUTE ON atype TO OBJTEST2;
CONNECT OBJTEST2/OBJTEST2;
CREATE TABLE atab OF OBJTEST1.atype;

BEGIN
  OBJTEST1.atype.newa(1, 'atab', 'OBJTEST2');
END;
```

/

The statement is successful this time because the procedure is executed under the privileges of the invoker (OBJTEST2), not the definer (OBJTEST1).

In a type hierarchy, a subtype has the same rights model as its immediate supertype. That is, it implicitly inherits the rights model of the supertype and cannot explicitly specify one. Furthermore, if the supertype was declared with definer rights, the subtype must reside in the same schema as the supertype. These rules allow invoker-rights type hierarchies to span schemas. However, type hierarchies that use a definer-rights model must reside within a single schema. For example:

```
CREATE TYPE deftype1 AS OBJECT (...); --Definer-rights type
CREATE TYPE subtype1 UNDER deftype1 (...); --subtype in same schema as supertype
CREATE TYPE schema2.subtype2 UNDER deftype1 (...); --ERROR
CREATE TYPE invtype1 AUTHID CURRENT_USER AS OBJECT (...); --Invoker-rights type
CREATE TYPE schema2.subtype2 UNDER invtype1 (...); --LEGAL
```

## Replicating Object Tables and Columns

Object tables and object views can be replicated as materialized views. You can also replicate relational tables that contain columns of an object, collection, or REF type. Such materialized views are called object-relational materialized views.

All user-defined types required by an object-relational materialized view must exist at the materialized view site as well as at the master site. They must have the same object type IDs and versions at both sites.

### Replicating Columns of Object, Collection, or REF Type

To be updatable, a materialized view based on a table that contains an object column must select the column as an object in the query that defines the view: if the query selects only certain attributes of the column's object type, then the materialized view is read-only.

The view-definition query can also select columns of collection or REF type. REFS can be either primary-key based or have a system-generated key, and they can be either scoped or unscoped. Scoped REF columns can be rescoped to a different table at the site of the materialized view—for example, to a local materialized view of the master table instead of the original, remote table.

## Replicating Object Tables

A materialized view based on an object table is called an object materialized view. Such a materialized view is itself an object table. An object materialized view is created by adding the *OF type* keyword to the `CREATE MATERIALIZED VIEW` statement. For example:

```
CREATE MATERIALIZED VIEW customer OF cust_objtyp
AS SELECT * FROM Scott.Customer_objtab@dbs1;
```

As with an ordinary object table, each row of an object materialized view is an object instance, so the view-definition query that creates the materialized view must select entire objects from the master table: the query cannot select only a subset of the object type's attributes. For example, the following materialized view is not allowed:

```
CREATE MATERIALIZED VIEW customer OF cust_objtyp
AS SELECT CustNo FROM Scott.Customer_objtab@dbs1;
```

You can create an object-relational materialized view from an object table by omitting the *OF type* keyword, but such a view is read-only: you cannot create an updatable object-relational materialized view from an object table.

For example, the following `CREATE MATERIALIZED VIEW` statement creates a read-only object-relational materialized view of an object table. Even though the view-definition query selects all columns and attributes of the object type, it does not select them as attributes of an object, so the view created is object-relational and read-only:

```
CREATE MATERIALIZED VIEW customer
AS SELECT * FROM Scott.Customer_objtab@dbs1;
```

For both object-relational and object materialized views that are based on an object table, if the type of the master object table is not `FINAL`, the `FROM` clause in the materialized view definition query must include the `ONLY` keyword. For example:

```
CREATE MATERIALIZED VIEW customer OF cust_objtyp
AS SELECT CustNo FROM ONLY Scott.Customer_objtab@dbs1;
```

Otherwise, the `FROM` clause must omit the `ONLY` keyword.

**See Also:** *Oracle Database Advanced Replication* for more information on replicating object tables and columns

## Constraints on Objects

Oracle does not support constraints and defaults in type specifications. However, you can specify the constraints and defaults when creating the tables:

```
CREATE TYPE customer_type AS OBJECT(  
    cust_id INTEGER);  
/  
CREATE TYPE department_type AS OBJECT(  
    deptno INTEGER);  
/  
CREATE TABLE customer_tab OF customer_type (  
    cust_id default 1 NOT NULL);  
  
CREATE TABLE department_tab OF department_type (  
    deptno PRIMARY KEY);  
  
CREATE TABLE customer_tab1 (  
    cust customer_type DEFAULT customer_type(1)  
    CHECK (cust.cust_id IS NOT NULL),  
    some_other_column VARCHAR2(32));
```

## Considerations Related to Type Evolution

The following sections contain design considerations relating to type evolution.

### Pushing a Type Change Out to Clients

Once a type has evolved on the server side, all client applications using this type need to make the necessary changes to structures associated with the type. You can do this with OTT/J PUB. You also may need to make programmatic changes associated with the structural change. After making these changes, you must recompile your application and relink.

Types may be altered between releases of a third-party application. To inform client applications that they need to recompile to become compatible with the latest release of the third-party application, you can have the clients call a release-oriented compatibility initialization function. This function could take as input a string that tells it which release the client application is working with. If the release string mismatches with the latest version, an error is generated. The client application must then change the release string as part of the changes required to become compatible with the latest release.

For example:

```
FUNCTION compatibility_init(rel IN VARCHAR2, errmsg OUT VARCHAR2)
RETURN NUMBER;
```

where:

- `rel` is a release string that is chosen by the product, such as, 'Release 10.1'
- `errmsg` is any error message that may need to be returned
- The function returns 0 on success and a nonzero value on error

## Changing Default Constructors

When a type is altered, its default, system-defined constructors need to be changed in order (for example) to include newly added attributes in the parameter list. If you are using default constructors, you need to modify their invocations in your program in order for the calls to compile.

You can avoid having to modify constructor calls if you define your own constructor functions instead of using the system-defined default ones.

**See Also:** ["Advantages of User-Defined Constructors"](#) on page 7-20

## Altering the FINAL Property of a Type

When you alter a type `T1` from `FINAL` to `NOT FINAL`, any attribute of type `T1` in the client program changes from being an inlined structure to a pointer to `T1`. This means that you need to change the program to use dereferencing when this attribute is accessed.

Conversely, when you alter a type from `NOT FINAL` to `FINAL`, the attributes of that type change from being pointers to inlined structures.

For example, say that you have the types `T1(a int)` and `T2(b T1)`, where `T1`'s property is `FINAL`. The C/JAVA structure corresponding to `T2` is `T2(T1 b)`. But if you change `T1`'s property to `NOT FINAL`, then `T2`'s structure becomes `T2(T1 *b)`.

## Parallel Queries with Oracle Objects

Oracle lets you perform parallel queries with objects and objects synthesized in views, when you follow these rules:

- To make queries involving joins and sorts parallel (using the `ORDER BY`, `GROUP BY`, and `SET` operations), a `MAP` function is required. In the absence of a `MAP` function, the query automatically becomes serial.
- Parallel queries on nested tables are not supported. Even if there are parallel hints or parallel attributes for the table, the query is serial.
- Parallel DML and parallel DDL are not supported with objects. DML and DDL are always performed in serial.
- Parallel DML is not supported on views with `INSTEAD-OF` trigger. However, the individual statements within the trigger may be parallelized.

## Design Consideration Tips and Techniques

The following sections provide assorted tips on various aspects of working with Oracle object types.

### Deciding Whether to Evolve a Type or Create a Subtype Instead

As an application goes through its life cycle, the question often arises whether to change an existing user-defined type or to create a specialized subtype to meet new requirements. The answer depends on the nature of the new requirements and their context in the overall application semantics. Here are two examples:

#### Changing a Widely Used Base Type

Suppose that we have a user-defined type `address` with attributes `Street`, `State`, and `ZIP`:

```
CREATE TYPE address AS OBJECT (  
    Street VARCHAR2(80),  
    State  VARCHAR2(20),  
    ZIP    VARCHAR2(10));  
/
```

We later find that we need to extend the `address` type by adding a `Country` attribute to support addresses internationally. Is it better to create a subtype of `address` or to evolve the `address` type itself?

With a general base type that has been widely used throughout an application, it is better to implement the change using type evolution.

### Adding Specialization

Suppose that an existing type hierarchy of Graphic types (for example, curve, circle, square, text) needs to accommodate an additional variation, namely, Bezier curve. To support a new specialization of this sort that does not reflect a shortcoming of the base type, we should use inheritance and create a new subtype `BezierCurve` under the `Curve` type.

To sum up, the semantics of the required change dictates whether we should use type evolution or inheritance. For a change that is more general and affects the base type, use type evolution. For a more specialized change, implement the change using inheritance.

### How ANYDATA Differs from User-Defined Types

`ANYDATA` is an Oracle-supplied type that can hold instances of any Oracle datatype, whether built-in or user-defined. `ANYDATA` is a self-describing type and supports a reflection-like API that you can use to determine the shape of an instance.

While both inheritance, through the substitutability feature, and `ANYDATA` provide the polymorphic ability to store any of a set of possible instances in a placeholder, the two models give the capability two very different forms.

In the inheritance model, the polymorphic set of possible instances must form part of a single type hierarchy. A variable can potentially hold instances only of its defined type or of its subtypes. You can access attributes of the supertype and call methods defined in the supertype (and potentially overridden by the subtype). You can also test the specific type of an instance using the `IS OF` and the `TREAT` operators.

`ANYDATA` variables, however, can store heterogeneous instances. You cannot access attributes or call methods of the actual instance stored in an `ANYDATA` variable (unless you extract out the instance). You use the `ANYDATA` methods to discover and extract the type of the instance. `ANYDATA` is a very useful mechanism for parameter passing when the function/procedure does not care about the specific type of the parameter(s).

Inheritance provides better modeling, strong typing, specialization, and so on. Use `ANYDATA` when you simply want to be able to hold one of any number of possible instances that do not necessarily have anything in common.

### Polymorphic Views: An Alternative to an Object View Hierarchy

[Chapter 5](#) describes how to build up a view hierarchy from a set of object views each of which contains objects of a single type. Such a view hierarchy enables

queries on a view within the hierarchy to see a polymorphic set of objects contained by the queried view or its subviews.

As an alternative way to support such polymorphic queries, you can define an object view based on a query that returns a polymorphic set of objects. This approach is especially useful when you want to define a view over a set of tables or views that already exists.

For example, an object view of `Person_t` can be defined over a query that returns `Person_t` instances, including `Employee_t` instances. The following statement creates a view based on queries that select persons from a `persons` table and employees from an `employees` table.

```
CREATE VIEW Persons_view OF Person_t AS
  SELECT Person_t(...) FROM persons
  UNION ALL
  SELECT TREAT(Employee_t(...) AS Person_t) FROM employees;
```

An `INSTEAD OF` trigger defined for this view can use the `VALUE` function to access the current object and to take appropriate action based on the object's most specific type.

Polymorphic views and object view hierarchies have these important differences:

- **Addressability:** In a view hierarchy, each subview can be referenced independently in queries and DML statements. Thus, every set of objects of a particular type has a logical name. However, a polymorphic view is a single view, so you must use predicates to obtain the set of objects of a particular type.
- **Evolution:** If a new subtype is added, a subview can be added to a view hierarchy without changing existing view definitions. With a polymorphic view, the single view definition must be modified by adding another `UNION` branch.
- **DML Statements:** In a view hierarchy, each subview can be either inherently updatable or can have its own `INSTEAD OF` trigger. With a polymorphic view, only one `INSTEAD OF` trigger can be defined for a given operation on the view.

## The SQLJ Object Type

This section discusses the SQLJ object type.

### The Intended Use of SQLJ Object Types

According to the *Information Technology - SQLJ - Part 2* document (SQLJ Standard), a SQLJ object type is a database object type designed for Java. A SQLJ object type

maps to a Java class. Once the mapping is registered through the extended SQL `CREATE TYPE` command (a DDL statement), the Java application can insert or select the Java objects directly into or from the database through an Oracle JDBC driver. This enables the user to deploy the same class in the client, through JDBC, and in the server, through SQL method dispatch.

### Actions Performed When Creating a SQLJ Object Type

The extended SQL `CREATE TYPE` command:

- Populates the database catalog with the external names for attributes, functions, and the Java class. Also, dependencies between the Java class and its corresponding SQLJ object type are maintained.
- Validates the existence of the Java class and validates that it implements the interface corresponding to the value of the `USING` clause.
- Validates the existence of the Java fields (as specified in the `EXTERNAL NAME` clause) and whether these fields are compatible with corresponding SQL attributes.
- Generates an internal class to support constructors, external variable names, and external functions that return `self` as a result.

### Uses of SQLJ Object Types

The SQLJ object type is a special case of SQL object type in which all methods are implemented in a Java class(es). The mapping between a Java class and its corresponding SQL type is managed by the SQLJ object type specification. That is, the SQLJ Object type specification cannot have a corresponding type body specification.

Also, the inheritance rules among SQLJ object types specify the legal mapping between a Java class hierarchy and its corresponding SQLJ object type hierarchy. These rules ensure that the SQLJ Type hierarchy contains a valid mapping. That is, the supertype or subtype of a SQLJ object type has to be another SQLJ object type.

### Uses of Custom Object Types

The custom object type is the Java interface for accessing SQL object type. A SQL object type may include methods that are implemented in languages such as PLSQL, Java, and C. Methods implemented in Java in a given SQL object type can belong to different unrelated classes. That is, the SQL object type does not map to a specific Java class.

In order for the client to access these objects, JPub can be used to generate the corresponding Java class. Furthermore, the user has to augment the generated classes with the code of the corresponding methods. Alternatively, the user can create the class corresponding to the SQL object type.

At runtime, the JDBC user has to register the correspondence between a SQL Type name and its corresponding Java class in a map.

### Differences Between SQLJ and Custom Object Types Through JDBC

The following table summarizes the differences between SQLJ object types and custom object types.

**Table 8–1 Differences Between SQLJ and Custom Object Types**

Feature	SQLJ Object Type Behavior	Custom Object Type Behavior
Typecodes	Use the <code>OracleTypes.JAVA_STRUCT</code> typecode to register a SQLJ object type as a SQL OUT parameter. The <code>OracleTypes.JAVA_STRUCT</code> typecode is also used in the <code>_SQL_TYPECODE</code> field of a class implementing the <code>ORADData</code> or <code>SQLData</code> interface.	Use the <code>OracleTypes.STRUCT</code> typecode to register a custom object type as a SQL OUT parameter. The <code>OracleTypes.STRUCT</code> typecode is also used in the <code>_SQL_TYPECODE</code> field of a class implementing the <code>ORADData</code> or <code>SQLData</code> interface.
Creation	Create a Java class implementing the <code>SQLData</code> or <code>ORADData</code> and <code>ORADDataFactory</code> interfaces first and then load the Java class into the database. Next, you issue the extended SQL <code>CREATE TYPE</code> command for SQLJ object type.	Issue the extended SQL <code>CREATE TYPE</code> command for a custom object type and then create the <code>SQLData</code> or <code>ORADData</code> Java wrapper class using <code>JPublisher</code> or do this manually.
Method Support	Supports external names, constructor calls, and calls for member functions with side effects.	There is no default class for implementing type methods as Java methods. Some methods may also be implemented in SQL.
Type Mapping	Type mapping is automatically done by the extended SQL <code>CREATE TYPE</code> command. However, the SQLJ object type must have a defining Java class on the client.	Register the correspondence between SQL and Java in a type map. Otherwise, the type is materialized as <code>oracle.sql.STRUCT</code> .
Inheritance	There are rules for mapping SQL hierarchy to a Java class hierarchy. See the <i>Oracle Database SQL Reference</i> for a complete description of these rules.	There are no mapping rules.

## Miscellaneous Tips

This section discusses miscellaneous tips.

### Column Substitutability and the Number of Attributes in a Hierarchy

If a column or table is of type `T`, Oracle adds a hidden column for each attribute of type `T` and, if the column or table is substitutable, for each attribute of every

subtype of `T`, to store attribute data. A hidden `typeid` column is added as well, to keep track of the type of the object instance in a row.

The number of columns in a table is limited to 1,000. A type hierarchy with a number of total attributes approaching 1,000 puts you at risk of running up against this limit when using substitutable columns of a type in the hierarchy. To avoid problems as a result of this, consider one of the following options for dealing with a hierarchy that has a large number of total attributes:

- Use views
- Use `REFS`
- Break up the hierarchy

### **Circular Dependencies Among Types**

Avoid creating circular dependencies among types. In other words, do not create situations in which a method of type `T` returns a type `T1`, which has a method that returns a type `T`.

---

---

# A Sample Application Using Object-Relational Features

This chapter describes a sample application that provides an overview of how to create and use user-defined datatypes (Oracle Objects). An application is first developed with the relational model and then with the object-relational model.

---

---

**Note:** The extended sample application description is not included in the PDF version of this chapter but is included in the HTML version of the chapter.

---

---

This chapter contains the following sections:

- [Introduction to the Sample Application](#)

## Introduction to the Sample Application

User-defined types are schema objects in which users formalize the data structures and operations that appear in their applications.

The examples in this chapter illustrate the most important aspects of defining, using, and evolving user-defined types. One important aspect of working with user-defined types is creating methods that perform operations on objects. In the example, definitions of object type methods use the PL/SQL language. Other aspects of using user-defined types, such as defining a type, use SQL.

The examples develop different versions of a database schema for an application that manages customer purchase orders. First a purely relational version is shown, and then an equivalent, object-relational version. Both versions provide for the same basic kinds of entities—customers, purchase orders, line items, and so on. But the object-relational version creates user-defined types for these entities and manages data for particular customers and purchase orders by instantiating instances of the respective user-defined types.

PL/SQL and Java provide additional capabilities beyond those illustrated in this chapter, especially in the area of accessing and manipulating the elements of collections.

Client applications that use the Oracle Call Interface (OCI), Pro\*C/C++, or Oracle Objects for OLE (OO4O) can take advantage of its extensive facilities for accessing objects and collections, and manipulating them on clients.

### See Also:

- *Oracle Database SQL Reference* for a complete description of SQL syntax and usage for user-defined types
- *PL/SQL User's Guide and Reference* for a complete discussion of PL/SQL capabilities
- *Oracle Database Java Developer's Guide* for a complete discussion of Java
- *Oracle Call Interface Programmer's Guide*
- *Pro\*C/C++ Programmer's Guide*

## A

---

Active Server Pages, 4-10  
ActiveX, 4-10  
ADMIN OPTION  
    with EXECUTE ANY TYPE, 6-3  
aggregate functions  
    *See* user-defined aggregate functions  
ALTER ANY TYPE privilege, 6-2  
    *See also* privileges  
ALTER TABLE, 7-19  
    *See also* object types, evolving  
ALTER TYPE statement, 4-19, 7-15  
    *See also* object types, evolving  
ANYDATA datatype, 7-25, 8-36  
ANYDATASET datatype, 7-25  
ANYTYPE datatype, 7-25  
arrays  
    size of VARRAYs, 3-3  
    variable (VARRAYs), 3-3  
ASP, 4-10  
atomic nulls, 2-2  
attributes  
    leaf-level, 7-2  
    leaf-level scalar, 7-2  
    modifying, 7-13  
    of object types, 1-5

## B

---

bind variables  
    user-defined types, 4-2

## C

---

caches  
    object cache, 4-2, 4-7, 5-4, 6-5  
    object views, 5-4  
capture avoidance rule, 2-7  
CARDINALITY function, 3-19  
CAST function, 2-35  
COLLECT function, 3-20  
collections  
    assigning, 2-34  
    comparing, 3-11, 3-17  
    constructing, 1-12  
    creating, 3-2  
    datatypes  
    DML on, 3-15  
    multilevel, 3-9, 8-11  
        constructing, 3-11  
        creating, 3-11  
        creating with REFs, 8-19  
        DML, 3-16  
        object views containing, 5-7  
    nested tables, 3-4  
    querying, 3-12, 8-11  
    *See also* varrays, nested tables  
    substituting in, 2-24  
    supported datatypes, 1-12  
    variable arrays (VARRAYs), 3-3  
column objects  
    versus row objects, 8-2  
COLUMN\_VALUE keyword, 3-10  
columns  
    column objects, 1-9  
    hidden, 7-2, 7-5

- indexes on column objects, 2-5
- qualifying in queries, 2-6
- qualifying names in queries, 2-7
- comparisons
  - collections, 3-17
  - methods, 2-10
- compilation of object types, 6-6
- COMPRESS clause
  - nested tables, 8-16
- CONNECT role
  - user-defined types, 6-2, 6-3
- constraints
  - object tables, 2-4
  - on Oracle objects, 8-33
  - REFs, 8-8
- constructor methods, 1-7, 2-14, 7-2
  - literal invocation of, 2-14
- constructors
  - methods, 2-14
  - user-defined, 7-20
- CREATE ANY TYPE privilege, 6-2
  - See also* privileges
- CREATE INDEX statement
  - object types, 2-5
- CREATE TABLE statement
  - examples
    - column objects, 1-6
    - object tables, 1-8, 2-4
- CREATE TRIGGER statement
  - examples
    - object tables, 2-5
- CREATE TYPE privilege, 6-2
  - See also* privileges
- CREATE TYPE statement, 8-39
  - collection types, 1-12
  - example, 3-5
  - incomplete types, 6-6
  - nested tables, 2-14, 3-5
  - object types, 1-4, 1-6, 2-2, 2-14
  - varrays, 3-4
- creating VARRAYs
  - containing references to LOBs, 3-8
- CURSOR expression, 2-36

## D

---

- dangling REFs, 1-11
- database administrators (DBAs)
  - DBA role, 6-2
- database links, 2-8
- datatypes
  - array types, 3-3
  - nested tables, 3-4
  - See also* object types, user-defined types
  - transient and generic, 7-25
- DBA role
  - user-defined types, 6-2
- default values
  - collections, 3-3
  - user-defined types, 3-3
- DELETE privilege for object tables, 6-4, 6-5
- DEREF function, 2-36
- dereferencing, 1-11
  - implicit, 1-11
- dot notation, 2-10
- DROP ANY TYPE privilege, 6-2
  - See also* privileges
- DROP TYPE statement, 6-9
- dump files
  - Export and Import, 6-14

## E

---

- equal and not equal conditions
  - nested tables, 3-17
- Excel, 4-10
- EXECUTE ANY TYPE privilege, 6-2, 6-3
  - See also* privileges
- EXECUTE privilege
  - user-defined types, 6-3
  - See also* privileges
- Export utility
  - user-defined types, 6-14
- EXTERNAL NAME phrase, 4-16

## F

---

- files
  - Export and Import dump file, 6-14
- FINAL keyword, 2-17

- modifying finality, 7-14, 8-34
- FORCE keyword, 5-18
- FORCE option, 6-9
- function-based indexes
  - on type methods, 8-28

## G

---

- GRANT option for EXECUTE privilege, 6-3
- granting
  - execute user-defined type, 6-3

## I

---

- implicit dereferencing, 1-11
- Import utility
  - user-defined types, 6-14
- IN condition, 3-18
- incomplete object types, 6-6
- indexes
  - on REFS, 2-5
  - type-discriminant column, 7-5
  - user-defined types, 2-5
- index-organized tables
  - storing nested tables as, 3-11, 8-15
- inheritance, 1-7
  - See type inheritance
- inner capture, 2-7
- INSERT privilege for object tables, 6-4, 6-5
- INSTANTIABLE keyword, 2-19
  - modifying instantiability, 7-14
- INSTEAD OF triggers
  - nested tables, 5-13
- invoker-rights
  - object types, 8-29
- IS A SET condition, 3-19
- IS EMPTY condition, 3-19
- IS OF *type* predicate, 2-36

## J

---

- Java
  - object storage, 4-16
  - Oracle JDBC and Oracle objects, 4-13
  - Oracle SQLJ and Oracle objects, 4-14

- with Oracle objects, 4-12

JDBC  
See Oracle JDBC

JPublisher, 4-15

## L

---

- leaf-level attributes, 7-2
  - scalar, 7-2
- literal invocation
  - constructor methods, 2-14
- locators
  - returning nested tables as, 7-30, 8-18
  - using a hint, 8-19
- locks
  - object level locking, 4-3

## M

---

- managing
  - object types, 6-1
- map methods, 2-10, 8-7
  - comparing collections, 3-17
  - for comparing objects, 2-10
- materialized views, 8-32
- MEMBER condition, 3-18
- methods, 1-6, 2-8
  - choosing a language for, 8-25
  - comparison methods, 2-10
    - in a type hierarchy, 2-13
  - constructor methods, 2-14, 7-2
    - literal invocation, 2-14
  - dropping, 7-13
  - dynamic method dispatch, 2-23
  - execution privilege for, 6-2
  - final, 2-17
  - function-based indexes, 8-28
  - inheriting, 2-20
  - instantiability, 2-19
  - invoking, 2-10
  - map, 2-10, 8-7
  - map for comparing objects, 2-10
  - map required for collections, 3-17
  - member, 2-9
  - of object types, 1-5

- order, 2-12, 8-7
- overloading, 2-20
- overriding, 2-18, 2-20, 2-21
- PL/SQL, 4-2
- SELF parameter, 2-9
- static, 2-13, 8-27
- multilevel collections
  - See collections, multilevel
- MULTISET EXCEPT operator, 3-20
- MULTISET INTERSECT operator, 3-20
- multiset operations
  - with nested tables, 3-19
- MULTISET UNION operator, 3-21

## N

---

- narrowing, 2-33, 2-40
- nested tables, 3-4, 8-14
  - COMPRESS clause, 8-16
  - creating, 3-2
  - creating indexes on, 8-17
  - equal and not equal conditions, 3-17
  - in an index-organized table, 3-11, 8-15
  - indexes, 2-5
  - INSTEAD OF triggers, 5-13
  - multiset operations, 3-19
  - querying, 3-12
    - unnesting results, 3-13
  - returning as locators, 8-18
  - specifying storage in a tablespace, 3-6
  - storage, 3-9, 8-14
  - updating in views, 5-13
- NESTED\_TABLE\_GET\_REFS hint, 8-19
- NESTED\_TABLE\_ID keyword, 3-11, 8-17
- NLS\_LENGTH\_SEMANTICS initialization
  - parameter, 2-3
- NOCOPY compiler hint, 2-10
  - performance issues, 8-28
- NOT MEMBER condition, 3-18
- nulls
  - atomic, 2-2
  - object types, 2-2

## O

---

- object cache
  - object views, 5-4
  - OCI, 4-2
  - privileges, 6-5
  - Pro\*C, 4-7
- object identifiers, 1-8, 8-6
  - column and index, 8-6
  - for object types, 7-2
  - primary-key based, 8-6
  - REFs, 8-6
  - storage, 8-6
- object instances, 1-4, 1-5
  - comparing, 2-34
- object tables, 1-8, 8-5
  - constraints, 2-4
  - indexes, 2-5
  - replicating, 8-31
  - row objects, 1-9
  - triggers, 2-5
  - virtual object tables, 5-2
- object types, 1-2
  - assignments across, 2-32
  - attributes of, 1-5
  - column objects, 1-9
  - column objects versus row objects, 8-2
  - comparison methods for, 2-10
  - constructor methods for, 2-14, 7-2
  - creating, 1-6
  - creating subtypes of, 2-19
  - dependents, 6-6, 7-7
  - evolving, 7-7, 8-35
    - design considerations, 8-33
    - SQLJ types, 4-19
  - FINAL or NOT FINAL, 2-17
  - final/not final, 8-34
  - incomplete, 6-6, 6-7
  - indexes on column objects, 2-5
  - indexing, 7-5
  - instantiable/not instantiable, 2-19
  - invoker-rights, 8-29
  - locking in cache, 4-3
  - managing, 6-1
  - methods of, 1-5

- PL/SQL, 4-2
- mutually dependent, 6-6
- Oracle type translator, 4-8
- remote access to, 2-8, 5-14
- row objects, 1-9
- See also* type inheritance
- specializing, 1-7
- SQLJ types, 4-16
- substituting, 2-24
- use of table aliases, 2-7
- object views, 5-1 to 5-19
  - advantages of, 5-2
  - circular references, 5-16
  - defining, 5-3
  - hierarchies, 5-19, 8-36
    - privileges, 5-27
    - querying in, 5-26
  - modeling relationships, 5-11, 5-15
  - multilevel collections in, 5-7
  - nested tables, 5-13
  - null objects in, 5-6
  - OIDs with, 5-9
  - REFs to, 5-10
  - replicating, 8-31
  - updating through INSTEAD OF triggers, 5-12
- OBJECT\_ID pseudocolumn, 2-26
- OBJECT\_VALUE pseudocolumn, 2-26
- object-relational model, 9-1
  - comparing objects, 8-7
  - constraints, 8-33
  - design considerations, 8-1
  - limitations of relational model, 1-2
  - methods, 1-6, 2-8
  - partitioning, 7-29
  - programmatically environments for, 4-1, 4-12
  - replication, 8-31
- objects
  - collection objects, 5-6
  - comparing, 2-34
  - in columns, 5-4
  - object references, 5-10
  - row objects and object identifiers, 5-6
- OCCI, 4-9
- OCI
  - associative access, 4-3

- for Oracle objects
  - building a program, 4-5
  - navigational access, 4-4
  - object cache, 4-4
  - OCIObjectFlush, 5-4
  - OCIObjectPin, 5-4
- OIDs, 1-8
- Oracle JDBC
  - accessing Oracle object data, 4-13
- Oracle objects
  - See* object-relational model
- Oracle Objects for OLE
  - OraCollection interface, 4-12
  - OraObject interface, 4-11
  - OraRef interface, 4-12
  - support of object-relational features, 4-10
- Oracle SQLJ
  - creating custom Java classes, 4-14
  - JPublisher, 4-14
  - support for Oracle objects, 4-14
- Oracle type translator (OTT), 4-8
- OraCollection interface, 4-12
- ORADATA interface, 4-17
- OraObject interface, 4-11
- OraRef interface, 4-12
- order methods, 2-12, 8-7
- OTT, 4-8

## P

---

- parallel query
  - objects, 8-34
  - restrictions for Oracle objects, 8-34
  - view objects, 8-34
- partitioning
  - tables containing Oracle objects, 7-29
- pkREFs, 7-4
- PL/SQL
  - bind variables
    - user-defined types, 4-2
  - object views, 5-4
  - polymorphism, 2-15, 8-36
  - See also* substitutability
- POWERMULTISET function, 3-22
- POWERMULTISET\_BY\_CARDINALITY

- function, 3-22
- primary-key-based REFS, 7-4
- privileges
  - system
    - user-defined types, 6-2
  - user-defined types
    - acquired by role, 6-2
    - ALTER ANY TYPE, 6-2
    - checked when pinning, 6-5
    - column level for object tables, 6-5
    - CREATE ANY TYPE, 6-2
    - CREATE TYPE, 6-2
    - DELETE, 6-4, 6-5
    - DROP ANY TYPE, 6-2
    - EXECUTE, 6-3
    - EXECUTE ANY TYPE, 6-2, 6-3
    - EXECUTE ANY TYPE with ADMIN OPTION, 6-3
    - EXECUTE with GRANT option, 6-3
    - INSERT, 6-4, 6-5
    - SELECT, 6-4, 6-5
    - system privileges, 6-2
    - UPDATE, 6-4, 6-5
    - using, 6-3

#### Pro\*C/C++

- associative access, 4-7
- converting between Oracle and C types, 4-8
- navigational access, 4-7
- user-defined datatypes, 4-2
- programmatically environments
  - for Oracle objects, 4-1, 4-12

## Q

---

#### queries

- set membership, 8-19
- unnesting, 8-11
- varrays, 8-14

## R

---

- REF function, 2-38
- references *See* REFS
- REFs, 1-9
  - comparing, 2-35

- constraints on, 2-6, 8-8
- constructing from object identifiers, 7-2
- dangling, 1-11, 2-6
- dereferencing of, 1-11
- for rows of object views, 5-3
- implicit dereferencing of, 1-11
- indexes on, 2-5
- indexing, 8-9
- obtaining, 1-11
- pinning, 5-4, 6-5
- scoped, 1-10, 2-6, 7-4, 8-9
- size of, 7-4
- storage, 8-8
- substitutability and, 2-24, 2-28
- use of table aliases, 2-8
- WITH ROWID option, 8-10
- RESOURCE role
  - user-defined types, 6-2, 6-3
- roles
  - CONNECT role, 6-2, 6-3
  - DBA role, 6-2
  - RESOURCE role, 6-2, 6-3
- row objects, 1-9
  - storage, 8-5

## S

---

#### schemas

- qualifying column names, 2-7
- user-defined datatypes, 4-2
- user-defined types, 1-5
- scoped REFS, 1-10, 7-4
- SELECT privilege for object tables, 6-4, 6-5
- SELF parameter, 2-9
- SET function, 3-22
- SQL
  - user-defined datatypes, 4-2
    - embedded SQL, 4-7
    - OCI, 4-3
- SQLData interface, 4-17
- SQLJ
  - See* Oracle SQL
- SQLJ object types, 4-12, 8-37
  - creating, 4-17
  - mapping Java classes, 4-18

*See also* object types, Oracle SQLJ

## storage

- column objects, 8-2
- nested tables, 7-4
- object tables, 7-2
- REFs, 7-4

SUBMULTISET condition, 3-18

substitutability, 2-24

- attribute, 2-24
- collections and, 2-24
- column and row, 2-25, 7-5
- constraining, 2-30
- dependencies, 6-8
- narrowing, 2-33
- OBJECT\_ID, 2-26
- OBJECT\_VALUE, 2-26
- turning off, 2-29
- views and, 2-25, 8-36
- widening, 2-32

subtypes, 2-15, 2-27

- creating, 2-18
- indexing attributes of, 7-6
- specializing, 8-35

supertypes, 2-15, 2-27

SYS\_TYPEID function, 2-38, 7-5

SYS.ANYDATA, 7-25

SYS.ANYDATASET, 7-25

SYS.ANYTYPE, 7-25

system privileges

- ADMIN OPTION, 6-3
- user-defined types, 6-2
- See also* privileges

## T

---

### TABLE

- function, 2-39

TABLE expressions, 3-13, 8-11

### tables

- aliases, 2-7
- functions, 2-39
- nested tables, 3-4
  - indexes, 2-5
- object tables, 1-8
  - constraints, 2-4

indexes, 2-5

triggers, 2-5

virtual, 5-2

qualifying column names, 2-6, 2-7

TREAT function, 2-26, 2-33, 2-37, 2-39, 7-6

### triggers

INSTEAD OF triggers

object views and, 5-12

user-defined types, 2-5

type dependencies, 6-8

type evolution

*See* object types

type hierarchies, 1-7, 2-15

methods in, 2-13

type inheritance, 2-15

finality, 2-17

instantiability, 2-19

*See* inheritance

specializing subtypes, 2-16

typeids, 2-38, 7-5

types

*See* datatypes, object types

## U

---

UNDER keyword, 2-18

unnesting queries, 8-11

UPDATE privilege for object tables, 6-4, 6-5

### updates

object views, 5-12

user-defined aggregate functions, 7-28

user-defined constructors, 7-20

user-defined datatypes

*See also* user-defined types

user-defined types, 1-2

and remote databases, 2-8

collections

nested tables, 3-4

variable arrays (VARRAYs), 3-3

example of privileges, 6-3

Export and Import, 6-14

incomplete types, 6-5

managing, 6-1

object types

use of table aliases, 2-7

- performance tuning, 6-12
- privileges, 6-2
- See also* object types
- See also* object-relational model
- storage, 7-2
- synonyms, 6-9
- tools, 6-13
- utilities, 6-14

USING clause, 4-16

## V

---

validation

- failure, 7-15
- object types, 7-11

VALUE function, 2-41

variables

- bind variables
  - user-defined types, 4-2
- object variables, 5-4

varrays, 3-3

- accessing, 8-14
- creating, 3-2
- creating VARRAYs, 3-8
- increasing the number of elements, 3-8
- querying, 8-14
- storage, 3-6, 8-13
- updating, 8-14

views

- See also* object views
- substitutability, 2-25
- updatability, 5-12

Visual Basic, 4-10

## W

---

widening

- and substitutability, 2-32