

**Oracle9i**

Heterogeneous Connectivity Administrator's Guide

Release 1 (9.0.1)

June 2001

Part No. A88789-01

**ORACLE®**

---

Oracle9i Heterogeneous Connectivity Administrator's Guide, Release 1 (9.0.1)

Part No. A88789-01

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

Primary Author: Ted Burroughs

Contributing Authors: Vira Goorah and Raghu Mani

Contributors: Jacco Draaijer, Kishan Peyetti, Sridhar Rajogopal, Paul Raveling, and Eric Voss

Graphics Designer: Valerie Moore

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

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

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

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

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

Oracle is a registered trademark, and Oracle Net Services, SQL\*Plus, Oracle Call Interface, Oracle Transparent Gateway, Oracle7, Oracle7 Server, Oracle8, Oracle8i, Oracle9i, PL/SQL, Pro\*C, Pro\*C/C++, and Enterprise Manager 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>
<b>1 Introduction</b>	
<b>The Heterogeneous Challenge .....</b>	<b>1-2</b>
<b>The Heterogeneous Services Module in the Oracle Database Server.....</b>	<b>1-2</b>
<b>Integrating Heterogeneous Services Into the Oracle Server .....</b>	<b>1-4</b>
<b>Benefits of Heterogeneous Services .....</b>	<b>1-5</b>
Remote Data Access .....	1-5
Elimination of Unnecessary Data Duplication.....	1-5
Heterogeneous Database Integration .....	1-6
Application Development and End User Tools .....	1-6
Two-Phase Commit and Multi-Site Transactions.....	1-6
Query Optimization .....	1-6
Error Mapping and Logging.....	1-7
Pass-Through Feature .....	1-7
<b>2 Oracle Transparent Gateways and Generic Connectivity</b>	
<b>Heterogeneous Connectivity Process Architecture .....</b>	<b>2-2</b>
<b>Heterogeneous Services Agents.....</b>	<b>2-2</b>
<b>Types of Heterogeneous Services Agents .....</b>	<b>2-3</b>
Oracle Transparent Gateways .....	2-3
Generic Connectivity.....	2-4

<b>Heterogeneous Services Components</b> .....	2-4
Transaction Service.....	2-4
SQL Service.....	2-5
<b>Configuring Heterogeneous Services</b> .....	2-5
Data Dictionary Translations .....	2-6
Initialization Parameters.....	2-6
Capabilities .....	2-6
<b>The Heterogeneous Services Data Dictionary</b> .....	2-6
Classes and Instances .....	2-7
Data Dictionary Views .....	2-8
<b>Gateway Process Flow</b> .....	2-8
Oracle Transparent Gateways for Non-Oracle Database Systems .....	2-10

### 3 Major Features

<b>SQL and PL/SQL Support</b> .....	3-2
<b>Heterogeneous Replication</b> .....	3-3
<b>Passthrough SQL</b> .....	3-5
Using the <code>DBMS_HS_PASSTHROUGH</code> package .....	3-5
Considering the Implications of Using Pass-Through SQL .....	3-6
Executing Pass-Through SQL Statements .....	3-6
Executing Non-Queries .....	3-7
Executing Queries.....	3-11
<b>Result Set Support</b> .....	3-14
Introduction .....	3-14
Result Set Support In Non-Oracle Systems:.....	3-14
Model 1.....	3-15
Model 2.....	3-15
Heterogeneous Services Support for Result Sets .....	3-15
Cursor mode.....	3-16
Sequential Mode .....	3-16
Code Examples:.....	3-17
OCI program fetching from result sets in cursor mode.....	3-18
OCI program fetching from result sets in sequential mode.....	3-19
PL/SQL program fetching from result sets in cursor mode .....	3-21
<b>Data Dictionary Translations</b> .....	3-24

<b>Examples</b> .....	3-25
<b>Date Time</b> .....	3-27
<b>Two Phase Commit Protocol</b> .....	3-28
<b>Piecewise Long</b> .....	3-28
<b>SQL*Plus Describe Command</b> .....	3-29
<b>Constraints on SQL in a Distributed Environment</b> .....	3-29
Resolving Remote and Heterogeneous References .....	3-29
Resolving Important Restrictions.....	3-30
Updates, Inserts and Deletes.....	3-34
<b>Using Index and Table Statistics</b> .....	3-35
<b>Other Optimizations</b> .....	3-36
Remote Join Optimization .....	3-37
<b>Optimizer Restrictions for non-Oracle Access</b> .....	3-38

## 4 Using the Gateway

<b>Setting Up Access to Non-Oracle Systems</b> .....	4-2
Step 1: Install the Heterogeneous Services Data Dictionary .....	4-2
Step 2: Set Up the Environment to Access Heterogeneous Services Agents .....	4-2
A Sample Entry for a Oracle Net Service Name .....	4-3
A Sample Listener Entry .....	4-3
Step 3: Create the Database Link to the Non-Oracle System .....	4-4
Step 4: Test the Connection .....	4-4
<b>Initialization Parameters</b> .....	4-6
<b>Optimizing Data Transfers Using Bulk Fetch</b> .....	4-8
Using OCI, an Oracle Precompiler, or Another Tool for Array Fetches .....	4-9
Controlling the Array Fetch Between Oracle Database Server and Agent .....	4-9
Controlling the Array Fetch Between Agent and Non-Oracle Server .....	4-9
Controlling the Reblocking of Array Fetches .....	4-9
<b>Registering Agents</b> .....	4-11
Enabling Agent Self-Registration .....	4-11
Using Agent Self-Registration to Avoid Configuration Mismatches .....	4-12
Understanding Agent Self-Registration.....	4-13
Specifying <b>HS_AUTOREGISTER</b> .....	4-14
Disabling Agent Self-Registration.....	4-15
<b>Oracle Database Server SQL Construct Processing</b> .....	4-16

<b>Using Synonyms .....</b>	<b>4-17</b>
Example of a Distributed Query .....	4-17
<b>Copying Data from the Oracle Database Server to the Non-Oracle Database System .....</b>	<b>4-19</b>
<b>Copying Data from the Non-Oracle Database System to the Oracle Database Server .....</b>	<b>4-21</b>
<b>Heterogeneous Services Data Dictionary Views .....</b>	<b>4-22</b>
Understanding the Types of Views.....	4-22
Understanding the Sources of Data Dictionary Information .....	4-23
Using the General Views .....	4-24
Using the Transaction Service Views.....	4-25
Using the SQL Service Views.....	4-26
Using Views for Capabilities and Translations.....	4-26
Using Views for Data Dictionary Translations .....	4-26
<b>Using the Heterogeneous Services Dynamic Performance Views .....</b>	<b>4-28</b>
Determining Which Agents Are Running on a Host .....	4-28
Determining the Open Heterogeneous Services Sessions .....	4-28
Determining the Heterogeneous Services Parameters.....	4-29

## 5 Using Multithreaded Agents

<b>Concepts.....</b>	<b>5-2</b>
The Challenge of Dedicated Agent Architecture .....	5-2
The Advantage of Multithreading .....	5-2
<b>Multithreaded Agent Architecture .....</b>	<b>5-4</b>
Overview .....	5-4
The Monitor Thread .....	5-6
Dispatcher Threads.....	5-7
Task Threads.....	5-7
<b>Multithreaded Agent Administration.....</b>	<b>5-8</b>
Overview.....	5-8
Single Command Mode Commands.....	5-8
Shell Mode Commands.....	5-9

## 6 Performance Tips

<b>Optimizing Heterogeneous Distributed SQL Statements.....</b>	<b>6-2</b>
<b>Using Gateways and Partition Views .....</b>	<b>6-2</b>
<b>Optimizing Performance of Distributed Queries.....</b>	<b>6-3</b>

Choose the best SQL statement .....	6-3
Use the cost-based optimizer.....	6-3
Use views.....	6-3

## 7 Generic Connectivity

<b>What Is Generic Connectivity?</b> .....	7-2
Types of Agents .....	7-2
Generic Connectivity Architecture .....	7-3
Oracle and Non-Oracle Systems on Separate Machines.....	7-3
Oracle and Non-Oracle Systems on the Same Machine .....	7-4
SQL Execution.....	7-6
Data Type Mapping .....	7-6
Generic Connectivity Restrictions.....	7-6
<b>Supported Oracle SQL Statements</b> .....	7-7
Functions Supported by Generic Connectivity .....	7-7
<b>Configuring Generic Connectivity Agents</b> .....	7-8
Creating the Initialization File .....	7-8
Editing the Initialization File .....	7-8
Setting Initialization Parameters for an ODBC-based Data Source .....	7-10
Setting Agent Parameters on Windows NT .....	7-10
Setting Agent Parameters on UNIX platforms.....	7-11
Setting Initialization Parameters for an OLE DB-based Data Source .....	7-12
<b>ODBC Connectivity Requirements</b> .....	7-13
<b>OLE DB (SQL) Connectivity Requirements</b> .....	7-15
<b>OLE DB (FS) Connectivity Requirements</b> .....	7-16
Data Source Properties.....	7-18

## A Heterogeneous Services Initialization Parameters

<b>HS_COMMIT_POINT_STRENGTH</b> .....	A-3
<b>HS_DB_DOMAIN</b> .....	A-3
<b>HS_DB_INTERNAL_NAME</b> .....	A-3
<b>HS_DB_NAME</b> .....	A-4
<b>HS_DESCRIBE_CACHE_HWM</b> .....	A-4
<b>HS_FDS_CONNECT_INFO</b> .....	A-4
<b>HS_FDS_SHAREABLE_NAME</b> .....	A-6

<b>HS_FDS_TRACE_LEVEL</b> .....	A-6
<b>HS_LANGUAGE</b> .....	A-6
Character sets .....	A-7
Language .....	A-7
Territory .....	A-7
<b>HS_LONG_PIECE_TRANSFER_SIZE</b> .....	A-8
<b>HS-NLS_DATE_FORMAT</b> .....	A-8
<b>HS-NLS_DATE_LANGUAGE</b> .....	A-8
<b>HS-NLS_NCHAR</b> .....	A-9
<b>HS-NLS_TIMESTAMP_FORMAT</b> .....	A-9
<b>HS-NLS_TIMESTAMP_TZ_FORMAT</b> .....	A-10
<b>HS_OPEN_CURSORS</b> .....	A-10
<b>HS_ROWID_CACHE_SIZE</b> .....	A-10
<b>HS_RPC_FETCH_REBLOCKING</b> .....	A-11
<b>HS_RPC_FETCH_SIZE</b> .....	A-11
<b>HS_TIME_ZONE</b> .....	A-12
<b>IFILE</b> .....	A-12

## **B Data Type Mapping**

<b>Mapping ANSI Data Types to Oracle Data Types Through an ODBC Interface</b> .....	B-2
<b>Mapping ANSI Data Types to Oracle Data Types Through an OLE DB Interface</b> .....	B-4

## **C DBMS\_HS\_PASSTHROUGH for Pass-Through SQL**

<b>Summary of Subprograms</b> .....	C-3
<b>BIND_VARIABLE</b> procedure .....	C-4
Syntax .....	C-4
Parameters .....	C-5
Exceptions .....	C-5
Pragmas .....	C-5
<b>BIND_VARIABLE_NCHAR</b> procedure .....	C-6
Syntax .....	C-6
Parameters .....	C-6
Exceptions .....	C-6
Pragmas .....	C-6
<b>BIND_VARIABLE_RAW</b> procedure .....	C-7



Syntax.....	C-7
Parameters.....	C-7
Exceptions.....	C-7
Pragmas .....	C-8
<b>BIND_OUT_VARIABLE</b> procedure .....	C-8
Syntax.....	C-8
Parameters.....	C-9
Exceptions.....	C-9
Pragmas .....	C-10
<b>BIND_OUT_VARIABLE_NCHAR</b> procedure.....	C-10
Syntax.....	C-10
Parameters.....	C-10
Exceptions.....	C-11
Pragmas .....	C-11
<b>BIND_OUT_VARIABLE_RAW</b> procedure .....	C-11
Syntax.....	C-11
Parameters.....	C-11
Exceptions.....	C-12
Pragmas .....	C-12
<b>BIND_INOUT_VARIABLE</b> procedure .....	C-12
Syntax.....	C-13
Parameters.....	C-13
Exceptions.....	C-14
Pragmas .....	C-14
<b>BIND_INOUT_VARIABLE_NCHAR</b> procedure.....	C-14
Syntax.....	C-14
Parameters.....	C-15
Exceptions.....	C-15
Pragmas .....	C-15
<b>BIND_INOUT_VARIABLE_RAW</b> procedure.....	C-16
Syntax.....	C-16
Parameters.....	C-16
Exceptions.....	C-17
Pragmas .....	C-17
<b>CLOSE_CURSOR</b> function .....	C-17

Syntax .....	C-17
Parameter .....	C-18
Exceptions .....	C-18
Pragmas .....	C-18
<b>EXECUTE_IMMEDIATE</b> function .....	C-18
Syntax .....	C-18
Parameter Description .....	C-18
Returns .....	C-19
Exceptions .....	C-19
Pragmas .....	C-19
<b>EXECUTE_NON_QUERY</b> function .....	C-19
Syntax .....	C-19
Parameter .....	C-20
Returns .....	C-20
Exceptions .....	C-20
Pragmas .....	C-20
<b>FETCH_ROW</b> function .....	C-20
Syntax .....	C-21
Parameters and Descriptions .....	C-21
Returns .....	C-21
Exceptions .....	C-21
Pragmas .....	C-21
<b>GET_VALUE</b> procedure .....	C-22
Syntax .....	C-22
Parameters .....	C-23
Exceptions .....	C-23
Pragmas .....	C-23
<b>GET_VALUE_NCHAR</b> procedure .....	C-24
Syntax .....	C-24
Parameters .....	C-24
Exceptions .....	C-25
Pragmas .....	C-25
<b>GET_VALUE_RAW</b> procedure .....	C-25
Syntax .....	C-25
Parameters .....	C-26

Exceptions.....	C-26
Pragmas .....	C-26
<b>OPEN_CURSOR</b> function.....	C-27
Syntax.....	C-27
Returns .....	C-27
Exceptions.....	C-27
Pragmas .....	C-27
<b>PARSE</b> procedure.....	C-28
Syntax.....	C-28
Parameters .....	C-28
Exceptions.....	C-28
Pragmas .....	C-28

## **D Data Dictionary Translation Support**

<b>Accessing the Non-Oracle Data Dictionary</b> .....	D-1
<b>Heterogeneous Services Data Dictionary Views</b> .....	D-2
<b>Supported Views and Tables</b> .....	D-5
<b>Data Dictionary Mapping</b> .....	D-7
Generic Connectivity Data Dictionary Descriptions.....	D-8
<b>ALL_CATALOG</b> .....	D-8
<b>ALL_COL_COMMENTS</b> .....	D-9
<b>ALL_CONS_COLUMNS</b> .....	D-9
<b>ALL_CONSTRAINTS</b> .....	D-9
<b>ALL_IND_COLUMNS</b> .....	D-10
<b>ALL_INDEXES</b> .....	D-10
<b>ALL_OBJECTS</b> .....	D-12
<b>ALL_TAB_COLUMNS</b> .....	D-13
<b>ALL_TAB_COMMENTS</b> .....	D-14
<b>ALL_TABLES</b> .....	D-14
<b>ALL_USERS</b> .....	D-16
<b>ALL_VIEWS</b> .....	D-16
<b>DICTIONARY</b> .....	D-17
<b>USER_CATALOG</b> .....	D-17
<b>USER_COL_COMMENTS</b> .....	D-17
<b>USER_CONS_COLUMNS</b> .....	D-17

USER_CONSTRAINTS.....	D-18
USER_IND_COLUMNS.....	D-18
USER_INDEXES .....	D-19
USER_OBJECTS .....	D-21
USER_TAB_COLUMNS.....	D-21
USER_TAB_COMMENTS .....	D-22
USER_TABLES.....	D-23
USER_USERS .....	D-24
USER_VIEWS .....	D-25

---

---

# Send Us Your Comments

**Oracle9i Heterogeneous Connectivity Administrator's Guide, Release 1 (9.0.1)**

**Part No. A88789-01**

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this document. 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?

If you find any errors or have any other suggestions for improvement, please indicate the document title and part number, 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 (optionally) electronic mail address.



---

# Preface

*Oracle9i* Heterogeneous Connectivity Administrator's Guide describes implementation issues for Oracle9i Heterogeneous Connectivity and introduces the tools and utilities available to assist you in implementing and using this feature.

This preface contains these topics:

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

## Audience

*Oracle9i* Heterogeneous Connectivity Administrator's Guide is intended for the following users:

- Database administrators who want to administer distributed database systems that involve the following:
  - Oracle to Oracle database links
  - Oracle to non-Oracle database links
- Regular Oracle database server users who want to make use of the Heterogeneous Services feature in the Oracle database server
- Readers who want an high-level understanding of this product and how it works.

To use this document, you should be familiar with the following information:

- Relational database concepts and basic database administration as described in *Oracle9i Database Concepts* and the *Oracle9i Database Administrator's Guide*.
- The operating system environment under which database administrators are running Oracle.

## Organization

This document contains the following chapters:

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

Heterogeneous Services, an integrated module within the Oracle9i database server, has been designed to access data in non-Oracle systems by means of either Oracle Transparent Gateways or generic connectivity. This chapter introduces you to Heterogeneous Services by describing the kinds of situations in which Heterogeneous Services is needed and by explaining how Heterogeneous Services fulfills this need.

### **Chapter 2, "Oracle Transparent Gateways and Generic Connectivity"**

You can access a non-Oracle database system either by Transparent Gateways or with Generic Connectivity. This chapter describes the architecture of Heterogeneous Services insofar as it relates to each of these means of accessing a non-Oracle system.



### **Chapter 3, "Major Features"**

This chapter describes the major features provided by Heterogeneous Services.

### **Chapter 4, "Using the Gateway"**

This chapter explains how to use Oracle Transparent Gateways.

### **Chapter 5, "Using Multithreaded Agents"**

This chapter explains what multithreaded agents are, how they contribute to the overall efficiency of a distributed database system, and how to administer multithreaded agents.

### **Chapter 6, "Performance Tips"**

This chapter explains how to optimize distributed SQL statements, how to use partition views with Oracle Transparent Gateways, and how to optimize the performance of distributed queries.

### **Chapter 7, "Generic Connectivity"**

This chapter describes the configuration and usage of generic connectivity agents.

### **Appendix A, "Heterogeneous Services Initialization Parameters"**

This appendix lists Heterogeneous Services initialization parameters and gives instructions how to set them.

### **Appendix B, "Data Type Mapping"**

The tables in this appendix show how Oracle maps ANSI datatypes through ODBC and OLE DB interfaces to supported Oracle datatypes when it is retrieving data from a non-Oracle system.

### **Appendix C, "DBMS\_HS\_PASSTHROUGH for Pass-Through SQL"**

The package, `DBMS_HS_PASSTHROUGH`, contains the procedures and functions for pass-through SQL of Heterogeneous Services. This appendix documents each of them.

### **Appendix D, "Data Dictionary Translation Support"**

This appendix documents data dictionary translation support. It explains how to access non-Oracle data dictionaries, lists Heterogeneous Services data dictionary views, describes how to use supported views and tables, and explains data dictionary mapping.

## Related Documentation

For more information, see these Oracle resources:

- *Oracle9i Database Concepts*
- *Oracle9i Database Administrator's Guide*
- *Oracle9i Database New Features*

In North America, printed documentation is available for sale in the Oracle Store at

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

Customers in Europe, the Middle East, and Africa (EMEA) can purchase documentation from

<http://www.oraclebookshop.com/>

Other customers can contact their Oracle representative to purchase printed documentation.

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://technet.oracle.com/membership/index.htm>

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://technet.oracle.com/docs/index.htm>

Many of the examples in this book use the sample schemas of the seed database, which is installed by default when you install Oracle. Refer to *Oracle9i Sample Schemas* for information on how these schemas were created and how you can use them yourself.

In North America, printed documentation is available for sale in the Oracle Store at

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

Customers in Europe, the Middle East, and Africa (EMEA) can purchase documentation from

<http://www.oraclebookshop.com/>

Other customers can contact their Oracle representative to purchase printed documentation.

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://technet.oracle.com/membership/index.htm>

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://technet.oracle.com/docs/index.htm>

## 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>Oracle9i Database Concepts</i> Ensure that the recovery catalog and target database <i>do not</i> reside on the same disk.
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.

Convention	Meaning	Example
<code>lowercase monospace (fixed-width font)</code>	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 <code>sqlplus</code> to open SQL*Plus.  The password is specified in the <code>orapwd</code> file.  Back up the datafiles and control files in the <code>/disk1/oracle/dbs</code> directory.  The <code>department_id</code> , <code>department_name</code> , and <code>location_id</code> columns are in the <code>hr.departments</code> table.  Set the <code>QUERY_REWRITE_ENABLED</code> initialization parameter to <code>true</code> .  Connect as <code>oe</code> user.  The <code>JRePUTil</code> class implements these methods.
<code>lowercase monospace (fixed-width font) italic</code>	Lowercase monospace italic font represents placeholders or variables.	You can specify the <i>parallel_clause</i> .  Run <code>Uold_release.SQL</code> 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.	<code>DECIMAL (digits [ , precision ])</code>
{ }	Braces enclose two or more items, one of which is required. Do not enter the braces.	<code>{ENABLE   DISABLE}</code>
	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.	<code>{ENABLE   DISABLE}</code> <code>[COMPRESS   NOCOMPRESS]</code>

Convention	Meaning	Example
...	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>	<pre>CREATE TABLE ... AS subquery;  SELECT col1, col2, ... , coln FROM employees;</pre>
.	Vertical ellipsis points indicate that we have omitted several lines of code not directly related to the example.	
Other notation	You must enter symbols other than brackets, braces, vertical bars, and ellipsis points as shown.	<pre>acctbal NUMBER(11,2); acct      CONSTANT NUMBER(4) := 3;</pre>
<i>Italics</i>	Italicized text indicates placeholders or variables for which you must supply particular values.	<pre>CONNECT SYSTEM/system_password DB_NAME = database_name</pre>
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.	<pre>SELECT last_name, employee_id FROM employees; SELECT * FROM USER_TABLES; DROP TABLE hr.employees;</pre>
lowercase	Lowercase typeface indicates programmatic elements that you supply. For example, lowercase indicates names of tables, columns, or files.  <b>Note:</b> Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	<pre>SELECT last_name, employee_id FROM employees; sqlplus hr/hr  CREATE USER mjones IDENTIFIED BY ty3MU9;</pre>

## Documentation Accessibility

Oracle's goal is to make our products, services, and supporting documentation accessible to the disabled community with good usability. 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/>

**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.**

---

# Introduction

Heterogeneous Services, an integrated module within the Oracle9i database server, has been designed to access data in non-Oracle systems by means of either Oracle Transparent Gateways or generic connectivity. This chapter introduces you to Heterogeneous Services by describing the kinds of situations in which Heterogeneous Services is needed and by explaining how Heterogeneous Services fulfills this need.

This chapter contains these topics:

- [The Heterogeneous Challenge](#)
- [The Heterogeneous Services Module in the Oracle Database Server](#)
- [Integrating Heterogeneous Services Into the Oracle Server](#)
- [Benefits of Heterogeneous Services](#)

**Note Also:** For more information, please consult your gateway documentation for individual gateways.

## The Heterogeneous Challenge

Heterogeneous access is a challenge that affects many organizations. Many run several different database systems. Each of these systems stores data and has a set of applications that runs against it. Consolidation of this data into one database system is often difficult. This is in large part due to the fact that many of the applications that run against one database may not have an equivalent that runs against another. Until such time as migration to one consolidated database system is made feasible, it is necessary for the various heterogeneous database systems to work together.

There are several problems to overcome before such inter operability becomes possible. The database systems can have different access interfaces, different data types, different capabilities, and different ways of handling error conditions. Even when one relational database is trying to access another relational database the differences are significant. In such a situation, the common features of the databases include data access through SQL, two phase commit, and similar data types. However, there are significant differences as well. SQL dialects can be different as can transaction semantics. There can be some data types in one database that do not exist in the other. The most significant area of difference is in the data dictionaries of the two databases. Most data dictionaries contain similar information but the information is structured for each in a completely different way.

There are several possible ways of overcoming this problem. In this book, we describe the approach that Oracle has taken.

---

---

**Note:** The term "non-Oracle system" refers to the following:

- Any system accessed by PL/SQL procedures written in C (that is, by external procedures)
  - Any system accessed through SQL (that is, by Oracle Transparent Gateways or generic connectivity)
  - Any system accessed procedurally (that is, by procedural gateways)
- 
- 

## The Heterogeneous Services Module in the Oracle Database Server

If a client program wishes to access or modify data at several Oracle databases, it can open up connections to each of them. This approach, however, has several drawbacks. If data from the databases has to be joined, then the client will have to contain logic that does that. If data integrity has to be guaranteed, then the client will have to contain transaction coordination logic. An alternative approach is for the client to connect to one Oracle database and shift the burden of joining data and



transaction coordination to that database. We call the database that the client program connects to the **local** database. We call any database other than this one a **remote** database. The client program can access objects at any of the remote databases using database links. The Oracle query processor will take care of the joins and its transaction engine will take care of the transaction coordination.

The approach that Oracle has taken to solving the heterogeneous connectivity problem is to allow a non-Oracle system to be one of the remote nodes in the above scenario. From the client's point of view the remote non-Oracle system will function like a remote Oracle system would. It appears to understand the same SQL dialect and to have the same data dictionary structure as an Oracle system. Access to a non-Oracle system in this manner is done through a module in the Oracle server called Heterogeneous Services. Using Heterogeneous Services with the client program can do the following

- Retrieve and modify data stored in a non-Oracle system using Oracle SQL dialect.
- Execute stored procedures at the non-Oracle system using Oracle PL/SQL calls.
- Issue these SQL statements or PL/SQL calls from either Oracle client applications like SQL\*Plus or Oracle programmatic interfaces like Pro\*C or OCI.

---

---

**Note:** Heterogeneous Services can also be used to call external routines written in C using PL/SQL calls. This aspect of Heterogeneous Services is not covered in this book. For more information on external procedures please see *Oracle9i SQL Reference* and *Oracle9i Application Developer's Guide - Fundamentals*.

---

---

The work done by the Heterogeneous Services module is, for the most part, completely transparent to the end user. With only a few exceptions, you do not need to do anything different to access a non-Oracle system than you would for accessing an Oracle system. The Heterogeneous Services Module is used as enabling technology for many of heterogeneous access products that Oracle Corporation designs and for features including Oracle Transparent Gateways and Generic Connectivity (both of which are discussed in detail in this book).

You generally implement Heterogeneous Services in one of the following ways:

- You use an Oracle Transparent Gateway with Heterogeneous Services to access a particular, commercially available, non-Oracle system for which that Oracle Transparent Gateway has been designed. (For example, you use the Oracle

Transparent Gateway for Sybase on Solaris to access a Sybase database system operating on a Sun Solaris platform.)

- You use generic connectivity within Oracle Heterogeneous Services to access non-Oracle databases through an ODBC or a OLE DB interface.

## Integrating Heterogeneous Services Into the Oracle Server

Much of the processing power of Oracle Transparent Gateways for Oracle7 and earlier versions of the sever has been integrated into Oracle8i and later versions of the Oracle database server as a module called **Heterogeneous Services**.

In the all versions of the Oracle server and Oracle Transparent Gateways up to Oracle7, much of the transaction processing code for the gateway was contained in the gateway itself. However, much of the same code also existed within the Oracle database server. Because of this redundancy, using a gateway placed an unnecessary demand on system resources.

An additional redundancy existed whenever you tried to use more than one gateway from the same database server. This was because each gateway contained large segments of code that were common to all the gateways. This meant that using more than one gateway at a time also placed an unnecessary demand on system resources.

The approach that Oracle has taken for Oracle8i and later versions of the Oracle database server has been to integrate all code that is redundant in either of these two ways into the Heterogeneous Services module of the Oracle database server. The advantage to this is that using gateways now requires less memory storage space and processing power than it did in Oracle7 and earlier releases. The result is a "thin" transparent gateway which functions based on the Heterogeneous Services module of the database with the following benefits.

## Benefits of Heterogeneous Services

This section describes the following additional features provided by the Heterogeneous Services module:

- [Remote Data Access](#)
- [Elimination of Unnecessary Data Duplication](#)
- [Heterogeneous Database Integration](#)
- [Application Development and End User Tools](#)
- [Two-Phase Commit and Multi-Site Transactions](#)
- [Query Optimization](#)
- [Error Mapping and Logging](#)
- [Pass-Through Feature](#)

### Remote Data Access

Remote data access provides distributed database system administrators with several benefits.

Applications can take advantage of Oracle client-server capability to connect to a remote server using Oracle Net. The remote server can then connect to the gateway using a database link. So, because the Oracle architecture enables network connections between each of the components, you have more options for locating your data.

Remote access also gives you access to data outside your local environment. With remote access, you can move application development onto cost-efficient workstations or microcomputers. Also, with remote access, your data sources are virtually unlimited. Remote access also enables you to choose the best environment for your users. For example, data might be located on a platform that supports only character-mode interfaces, but, with remote access, users can access the data from desktop platforms that support graphical user interfaces.

### Elimination of Unnecessary Data Duplication

An Oracle Transparent Gateway gives applications direct access to non-Oracle database system data. This consequently eliminates the need to upload and download large amounts of data to different locations. Reducing the need to upload and download large amounts of data has the further consequence of reducing the

risk for unsynchronized or inconsistent data. And, by reducing the need for data duplication, an Oracle transparent gateway reduces the disk storage needs across all of your systems.

### **Heterogeneous Database Integration**

The Oracle database server can accept a SQL statement that queries data stored in several different databases. The Oracle database server with the Heterogeneous Services module processes the SQL statement and passes the appropriate SQL directly to other Oracle databases and through gateways to non-Oracle databases. The Oracle database server then combines the results and returns them to the client. This enables a query to be processed so that it spans the non-Oracle database system, other databases, and local and remote Oracle data.

### **Application Development and End User Tools**

An Oracle Transparent Gateway extends the range of user tools and application development that you can use to access the databases. These user tools increase application development and user productivity by reducing prototype, development, and maintenance time. This means that current Oracle users do not have to learn a new set of tools to access data stored in non-Oracle database system databases. Instead, they can access Oracle and non-Oracle database system data with a single set of tools. These tools can run on remote machines connected through Oracle Net to the Oracle database server.

### **Two-Phase Commit and Multi-Site Transactions**

In a distributed database system, the network might fail during a distributed transaction, raising the risk of data inconsistencies. The Oracle transaction model uses a two-phase commit protocol to protect the databases as the data is being committed at sites participating in a distributed transaction. This feature ensures that all database servers participating in the transaction must commit or roll back the transaction statements. The two-phase commit protocol is also supported (with some limitations) for non-Oracle systems when the user is accessing them through an Oracle Transparent Gateway.

### **Query Optimization**

Whenever possible, the Oracle database server passes the entire query to the non-Oracle system to utilize the indexes and statistics of the non-Oracle system tables.

When a query that involves multiple databases is processed, the Oracle database server passes optimized statements to the remote servers and gateways involved in the query to minimize the amount of data returned across the network.

## **Error Mapping and Logging**

The gateway provides error mapping and logging. It does this by mapping the non-Oracle database system error to an Oracle database server error message and adding all of the relevant error messages generated by non-Oracle database system. You can route messages to the client application, an operator console, an error log, or any combination of these destinations as needed. Error mapping provides database transparency for applications.

## **Pass-Through Feature**

As mentioned in the previous sections, Heterogeneous Services technology can allow clients to transparently access non-Oracle systems using Oracle SQL. In some cases, however, it becomes necessary to use non-Oracle system SQL to access the non-Oracle system. For such cases, Heterogeneous Services has a pass-through feature which allows the user to bypass Oracle's query processor and to issue non-Oracle system SQL to the non-Oracle system through the gateway.



---

# Oracle Transparent Gateways and Generic Connectivity

You can access a non-Oracle database system either by Transparent Gateways or with Generic Connectivity. This chapter describes the architecture of Heterogeneous Services insofar as it relates to each of these means of accessing a non-Oracle system.


This chapter contains these topics:

- [Heterogeneous Connectivity Process Architecture](#)
- [Heterogeneous Services Agents](#)
- [Types of Heterogeneous Services Agents](#)
- [Heterogeneous Services Components](#)
- [Configuring Heterogeneous Services](#)
- [The Heterogeneous Services Data Dictionary](#)
- [Gateway Process Flow](#)

## Heterogeneous Connectivity Process Architecture

At a high level, Oracle heterogeneous connectivity process architecture is structured as shown in [Figure 2-1](#).

**Figure 2-1 Oracle Heterogeneous Connectivity Process Architecture**



The Heterogeneous Services module in the kernel talks to a Heterogeneous Services agent process which, in turn, talks to the non-Oracle system. We can conceptually divide the code into three parts:

- The Heterogeneous Services Module in the Oracle database server. Most of the heterogeneous connectivity related processing is done in this module.
- Agent generic code. This is code in the agent that is generic to all Heterogeneous Services based products. This consists, for the most part, of code to communicate with the database and multithreading support.
- The driver. This is the module that communicates with the non-Oracle system. It is used to map calls from the Heterogeneous Services external application programming interface (API) onto the native API of the non-Oracle system and it is non-Oracle system specific.

## Heterogeneous Services Agents

An agent is the process through which an Oracle server connects to a non-Oracle system. The agent process consists of two components. These are agent generic code and a non-Oracle system-specific driver. An agent exists primarily to isolate the Oracle database server from third-party code. In order for a process to access the non-Oracle system, the non-Oracle system client libraries have to be linked into it. In the absence of the agent process, these libraries would have to be directly linked into the Oracle database and problems in this code could cause the Oracle server to go down. Having an agent process isolates the Oracle server from any problems in



third-party code so that even if a fatal error takes place, only the agent process will end.

An agent can reside in the following places:

- On the same machine as the non-Oracle system
- On the same machine as the Oracle server
- On a machine different from either of these two

Agent processes are usually started when a user session makes its first non-Oracle system access through a database link. These connections are made using Oracle's remote data access software, Oracle Net Services, which enables both client-server and server-server communication. The agent process continues to run until the user session is disconnected or the database link is explicitly closed.

Multithreaded agents behave slightly differently. They have to be explicitly started and shut down by a database administrator instead of automatically being spawned by Oracle Net Services.

**See Also:** For more information on multithreaded agents, please see [Chapter 5, "Using Multithreaded Agents"](#)

## Types of Heterogeneous Services Agents

### Oracle Transparent Gateways

An agent process that accesses a non-Oracle system is called a gateway. (Note that agents can also be used to execute external procedures.) Access to all gateways goes through the Heterogeneous Services module in the Oracle server and all gateways contain the same agent-generic code. Each gateway has a different driver linked in which maps the Heterogeneous Services application programming interface (API) to the client API of the non-Oracle system.

An Oracle Transparent gateway is a gateway that is designed for accessing a specific non-Oracle system. Oracle Corporation provides gateways to access several commercially produced non-Oracle systems; many of these gateways have been ported to several platforms. For example, an Oracle Transparent Gateway for Sybase on Solaris is the Solaris port of a gateway designed to access Sybase database systems.

With Oracle Transparent Gateways, you can use an Oracle database server to access data anywhere in a distributed database system without needing to know either the location of the data or how it is stored. When the results of your queries are

returned to you by the Oracle database server, they are presented to you as if the data stores from which they were taken all resided within a remote instance of an Oracle distributed database.

### Generic Connectivity

In addition to transparent gateways to various non-Oracle database systems, there is a set of agents that comprise the Oracle generic connectivity feature. These agents contain only generic code and the customer is responsible for providing the necessary drivers. Oracle has generic connectivity agents for ODBC and OLE DB that enable you to use ODBC and OLEDB drivers to access non-Oracle systems that have an ODBC or an OLE DB interface.

To build a gateway to a specific non-Oracle system using generic connectivity, you must connect an ODBC or OLE DB driver to the gateway for that non-Oracle system. These drivers are not provided by Oracle corporation. However, as long as Oracle Corporation supports the ODBC and OLE DB protocols, you can use the generic connectivity feature to access any non-Oracle system that can be accessed using an ODBC or OLE DB driver.

Generic connectivity has some limitations. The ODBC and OLEDB gateways have to be installed in the same Oracle Home directory as the Oracle database server. Connecting to one of these gateways from another Oracle database server is not supported. Functionality of these gateways, especially when compared to Oracle Transparent Gateways, is limited.

**See Also:** For more information, see [Chapter 7, "Generic Connectivity"](#)

## Heterogeneous Services Components

### Transaction Service

The transaction service component of the Heterogeneous Services module makes it possible for non-Oracle systems to be integrated into Oracle database server transactions and sessions. When you access a non-Oracle system for the first time over a database link within your Oracle user session, you transparently set up an authenticated session in the non-Oracle system. At the end of your Oracle user session, the authenticated session in the non-Oracle database system transparently closes at the non-Oracle system.

Additionally, one or more non-Oracle systems can participate in an Oracle distributed transaction. When an application commits a transaction, Oracle's two-phase commit protocol accesses the non-Oracle database system to coordinate transparently the distributed transaction. Even in those cases where the non-Oracle system does not support all aspects of Oracle two-phase commit protocol, Oracle can (with some limitations) support distributed transactions with the non-Oracle system.

## SQL Service

The standard query language (SQL) service handles the processing of all SQL-related operations. The work done by the SQL service includes:

1. Mapping Oracle internal SQL-related calls to the Heterogeneous Services driver application programming interface (API); this is in turn mapped by the driver to the client API of the non-Oracle system.
2. Translating SQL statements from Oracle's SQL dialect to the SQL dialect of the non-Oracle system.
3. Translating queries that reference Oracle data dictionary tables to queries that extract the necessary information from the non-Oracle system data dictionary.
4. Translating data from non-Oracle system data types to Oracle data types and back.
5. Making up for missing functionality at the non-Oracle system by issuing multiple queries to get the necessary data and doing post processing to get the desired results

## Configuring Heterogeneous Services

In the previous section, we described what the different heterogeneous components do. These components consist entirely of generic code and, in order to work with so many different non-Oracle systems, their behavior has to be configured. Each gateway has configuration information stored in the driver module and this information is uploaded to the server immediately after the connection to the gateway has been established. We can divide this configuration information into three parts:

- [Data Dictionary Translations](#)
- [Initialization Parameters](#)
- [Capabilities](#)

## Data Dictionary Translations

Data dictionary translations are views on non-Oracle system data dictionary tables that help Heterogeneous Services translate references to Oracle data dictionary tables into queries needed to retrieve the equivalent information from the non-Oracle system data dictionary.

---

---

**Note:** For a more detailed explanation of data dictionary translations, please see [Appendix D, "Data Dictionary Translation Support"](#).

---

---

## Initialization Parameters

Initialization parameters serve two functions.

- They give the user a means of fine-tuning the gateway to optimize performance and memory utilization for the gateway and the Heterogeneous Services module.
- They enable the user to tell the gateway (and, thereby, Heterogeneous Services) how the non-Oracle system has been configured (for example what language the non-Oracle system is running in). To put it another way, they give Heterogeneous Services information about the configurable properties of the non-Oracle system.

You can examine the initialization parameters for a session by querying the view `V$HS_PARAMETER`. Users can set initialization parameters in gateway initialization files.

## Capabilities

Capabilities tell Heterogeneous Services about the limitations of the non-Oracle system (such as what types of SQL statements are and are not supported) and how to map Oracle data types and SQL expressions to their non-Oracle system equivalents. In other words, they tell Heterogeneous Services about the non-configurable properties of the non-Oracle system. Capabilities cannot be changed by the user.

## The Heterogeneous Services Data Dictionary


As mentioned in the previous section, configuration information is uploaded from an agent to the Heterogeneous Services module immediately after the connection to

the agent has been established. Since this information can be very large in size, it is inefficient to do uploads on each connection. Therefore, the first time an Oracle database talks to an agent, the configuration information is uploaded and stored in Heterogeneous Services data dictionary tables and thereafter no upload takes place until something at the agent changes (for example, if a patch is applied or if the agent is upgraded to a new version).

## Classes and Instances

Using Heterogeneous Services, a user can access several non-Oracle systems from a single Oracle database. This is illustrated in [Figure 2-2](#)

**Figure 2-2 Accessing Multiple Non-Oracle Instances**



Both the agents upload configuration information that is stored as part of the Oracle data. This information is organized in the Heterogeneous Services data dictionary as follows.

In the Heterogeneous Services data dictionary, Oracle organizes data by two levels of granularity called class and instance. A class pertains to a specific type of non-Oracle system. For example, you might want to access the class of Sybase database systems with your Oracle database server. An instance defines specializations within a class. For example, you might want to access several separate instances within a Sybase database system. Instance information takes precedence over class information, and class information takes precedence over server-supplied defaults.

Although it is possible to store data dictionary information at one level of granularity by having completely separate definitions in the data dictionary for each individual instance, this might lead to an unnecessarily large amount of redundant data dictionary information. To avoid this, Oracle organizes the data dictionary by two levels of granularity, in which each class definition (one level of granularity) is shared by all the particular instances (a second level of granularity) under that class.

For example, suppose that the Oracle database server accesses three instances of Sybase and two instances of Ingres II. Sybase and Ingres II each have their own code, requiring separate class definitions for the Oracle database server to access them. The Heterogeneous Services data dictionary therefore would contain two class definitions, one for Sybase and one for Ingres II, with five instance definitions, one for each instance being accessed by the Oracle database server.

### Data Dictionary Views

The Heterogeneous Services data dictionary views contain the following kinds of information:

- Names of instances and classes uploaded into the Oracle data dictionary
- Capabilities, including SQL translations, defined for each class or instance
- Data Dictionary translations defined for each class or instance
- Initialization parameters defined for each class or instance

You can access information from the Oracle data dictionary by using fixed views. The views are categorized into three main types:


- General views
- Views used for the transaction service
- Views used for the SQL service

**See Also:** For more information on data dictionary views, see [Appendix D, "Data Dictionary Translation Support"](#)

### Gateway Process Flow

[Figure 2-3](#) shows a typical gateway process flow. The steps explain the sequence of events that occurs when a client application queries the non-Oracle database system database through the gateway.

Figure 2-3 Gateway Process Flow



1. The client application sends a query over Oracle Net to the Oracle database server.
2. The Oracle database server sends the query over to the gateway using Oracle Net.
3. For the first transaction in a session, the gateway logs into non-Oracle database system using a username and password that is valid in the non-Oracle system.
4. The gateway converts the Oracle SQL statement into a SQL statement understood by non-Oracle database system.
5. The gateway retrieves data using non-Oracle database system SQL statements.
6. The gateway converts retrieved data into a format compatible with the Oracle database server.
7. The gateway returns query results to the Oracle database server, again using Oracle Net Services.

8. The Oracle database server passes the query results to the client application by using Oracle Net. The database link remains open until the gateway session is finished or the database link is explicitly closed.

## Oracle Transparent Gateways for Non-Oracle Database Systems

Oracle client applications can access non-Oracle database system data with Oracle SQL just as if the data residing in the non-Oracle database system were stored in a remote Oracle database. Combined data residing in both Oracle and non-Oracle database system databases can be accessed by a single SQL statement performing heterogeneous joins and subselects. This means you can develop a single set of portable applications to use against both Oracle and non-Oracle database system databases. In this way, you can continue to develop new information systems without losing your investment in existing data and applications.

Also, transaction integrity for transactions involving updates to both Oracle and non-Oracle database system databases from a single Oracle database server is automatically protected by the Oracle two-phase commit feature.

Finally, synonyms in the Oracle database server can be used for transparent access to the non-Oracle system. Synonyms within the Oracle database server that point to database links to non-Oracle database system tables makes the physical location of the data transparent to the client application. This allows the future migration of data from the non-Oracle database system to Oracle to be transparent to client applications.

Only the Oracle database server and Oracle Net are needed to set up a gateway to a non-Oracle system. All other Oracle products are not necessary. However, using other Oracle products with the gateway can greatly extend the capabilities of a gateway.



---

## Major Features

This chapter describes the major features provided by Heterogeneous Services.

This chapter contains the following topics:

- SQL and PL/SQL Support
- Heterogeneous Replication
- Passthrough SQL
- Result Set Support
- Data Dictionary Translations
- Date Time
- Two Phase Commit Protocol
- Piecewise Long
- SQL\*Plus Describe Command
- Constraints on SQL in a Distributed Environment
- Using Index and Table Statistics
- Other Optimizations
- Optimizer Restrictions for non-Oracle Access

---

---

**Note:** Even though Heterogeneous Services has all these features, they are not necessarily available in all Heterogeneous Services based gateways. Not only must there be generic support for these features, which Heterogeneous Services provides, but there must also be support added to the driver for them. Please consult your gateways documentation to determine if a certain Heterogeneous Services feature is supported for your gateway.

---

---

## SQL and PL/SQL Support

SQL statements are translated and data types are mapped according to capabilities. PL/SQL calls are mapped to non-Oracle system stored procedures. In the case of SQL statements, if functionality is missing at the remote system, then either a simpler query is issued or the statement is broken up into multiple queries and the desired results are obtained by post processing in the Oracle database.

Even though Heterogeneous Services can, for the most part, incorporate non-Oracle systems into Oracle distributed sessions, there are several limitations to this. Some of the generic limitations are:

1. Data manipulation language statements that update objects on the remote non-Oracle system should not reference any objects on the local Oracle database. An example of such a statement is:

```
INSERT INTO remote_table@link as SELECT * FROM local_table;
```

Such statements will cause an error to be raised.

2. There is no support for `CONNECT BY` clauses in SQL statements.
3. ROWID support is limited; consult individual gateway documentation for more details. The Oracle Universal Rowid data type is not supported in any Oracle9i gateway.
4. LOBs, ADTs, and REFs are not supported.
5. PL/SQL in SQL is not supported. For example, a statement such as:

```
SELECT remote_func@link(a,b) FROM remote_table@link
```

will cause an error to be raised.

6. Remote packages are not supported.
7. Remote stored procedures can have `out` arguments of type ref cursor but not `in` or `in-out` objects.

8. None of the Oracle9i gateways supports shared database links.

---

---

**Note:** In addition to these generic limitation, each gateway can have additional limitations. Please consult the gateway documentation for individual gateways for a complete list of limitations of the product.

---

---

## Heterogeneous Replication

Data can be replicated between a non-Oracle system and an Oracle server with materialized views.

**See Also:** Oracle9i Replication for a full description of materialized views and replication facilities.

Materialized views instantiate data captured from tables at the non-Oracle master site at a particular point in time. This instant is defined by a refresh operation, which copies this data to the Oracle server and synchronizes the copy on Oracle with the master copy on the non-Oracle system. The "materialized" data is then available as a view on the Oracle server.

Replication facilities provide mechanisms to schedule refreshes and to collect materialized views into replication groups to facilitate their administration. Refresh groups permit refreshing multiple materialized views just as if they were a single object.

Heterogeneous replication support is necessarily limited to a subset of the full Oracle-to-Oracle replication functionality:

- Only the non-Oracle system can be the master site. This is because materialized views can be created only on an Oracle server.
- Materialized views must use complete refresh. This is because fast refresh would require Oracle-specific functionality in the non-Oracle system.
- Not all types of materialized views can be created to reference tables on a non-Oracle system. Primary key and subquery materialized views are supported, but rowid and object id materialized views are not supported. This is because there is no SQL standard for the format and contents of rowids, and non-Oracle systems do not implement Oracle objects.

Other restrictions apply to any access to non-Oracle data through Oracle's Heterogeneous Services facilities. The most important of these are:

- Non-Oracle data types in table columns mapped to a fixed view must be compatible with (that is, have a mapping to or from) Oracle data types. This is usually true for data types defined by ANSI SQL standards.
  - A subquery materialized view may not be able to use language features restricted by individual non-Oracle systems. In many cases Heterogeneous Services supports such language features by processing queries within the Oracle server, but occasionally the non-Oracle systems impose limitations that cannot be diagnosed until Heterogeneous Services attempts to execute the query.
9. The following examples illustrate basic setup and use of three materialized views to replicate data from a non-Oracle system to an Oracle data store.

---

---

**Note:** For the following examples, *remote\_db* refers to the non-Oracle system which you are accessing from your Oracle database server.

---

---

### Example 1: Set up 3 materialized views and a refresh group for them.

1. Create a primary key materialized view of table *thsmv\_customer@remote\_db*

```
CREATE MATERIALIZED VIEW pk_mv REFRESH COMPLETE AS
  SELECT * FROM thsmv_customer@remote_db WHERE "zip" = 94555;
```

2. Create a subquery materialized view of tables *thsmv\_orders@remote\_db* and *thsmv\_customer@remote\_db*

```
CREATE MATERIALIZED VIEW sq_mv REFRESH COMPLETE AS
  SELECT * FROM thsmv_orders@remote_db o WHERE EXISTS
    (SELECT c."c_id" FROM thsmv_customer@remote_db c
     WHERE c."zip" = 94555 and c."c_id" = o."c_id" );
```

3. Create a complex materialized view of data from multiple tables on *remote\_db*

```
CREATE MATERIALIZED VIEW cx_mv
  REFRESH COMPLETE AS
  SELECT c."c_id", o."o_id"
    FROM thsmv_customer@remote_db c,
         thsmv_orders@remote_db o,
         thsmv_order_line@remote_db ol
  WHERE c."c_id" = o."c_id"
     AND o."o_id" = ol."o_id";
```

### Example 2: Set up a refresh group for these 3 materialized views and force a refresh

```
BEGIN
  dbms_refresh.make('refgroup1',
    'pk_mv, sq_mv, cx_mv',
    NULL, NULL);
END;
/
```

### Example 3: Force refresh of all 3 materialized views

```
BEGIN
  dbms_refresh.refresh('refgroup1');
END;
/
```

## Passthrough SQL

The pass-through SQL feature allows you to send a statement directly to a non-Oracle system without being interpreted by the Oracle9i server. This feature can be useful if the non-Oracle system allows for operations in statements for which there is no equivalent in Oracle.

This section contains the following topics:

- [Passthrough SQL](#)
- [Considering the Implications of Using Pass-Through SQL](#)
- [Executing Pass-Through SQL Statements](#)

## Using the DBMS\_HS\_PASSTHROUGH package

You can execute Passthrough SQL statements directly at the non-Oracle system using the PL/SQL package DBMS\_HS\_PASSTHROUGH. Any statement executed with the pass-through package is executed in the same transaction as standard SQL statements.

The DBMS\_HS\_PASSTHROUGH package is a virtual package. It conceptually resides at the non-Oracle system. In reality, however, calls to this package are intercepted by Heterogeneous Services and mapped onto one or more Heterogeneous Services application programming interface (API) calls. The driver, in turn, maps these Heterogeneous Services API calls onto the API of the non-Oracle system. The client application should invoke the procedures in the package through a database link in exactly the same way as it would invoke a

non-Oracle system stored procedure. The special processing done by Heterogeneous Services is transparent to the user.

**See Also:** *Oracle9i Supplied PL/SQL Packages Reference* for more information about this package.

## Considering the Implications of Using Pass-Through SQL

When you execute a pass-through SQL statement that implicitly commits or rolls back a transaction in the non-Oracle system, the transaction is affected. For example, some systems implicitly commit the transaction containing a data definition language (DDL) statement. Because the Oracle database server is bypassed, the Oracle database server is unaware of the commit in the non-Oracle system. Consequently, the data at the non-Oracle system can be committed while the transaction in the Oracle database server is not.

If the transaction in the Oracle database server is rolled back, data inconsistencies between the Oracle database server and the non-Oracle server can occur. This situation results in **global data inconsistency**.

Note that if the application executes a regular COMMIT statement, the Oracle database server can coordinate the distributed transaction with the non-Oracle system. The statement executed with the pass-through facility is part of the distributed transaction.

## Executing Pass-Through SQL Statements

The table below shows the functions and procedures provided by the DBMS\_HS\_PASSTHROUGH package that allow you to execute pass-through SQL statements.

Procedure/Function	Description
OPEN_CURSOR	Opens a cursor
CLOSE_CURSOR	Closes a cursor
PARSE	Parses the statement
BIND_VARIABLE	Binds IN variables
BIND_OUT_VARIABLE	Binds OUT variables
BIND_INOUT_VARIABLE	Binds IN OUT variables
EXECUTE_NON_QUERY	Executes non-query
EXECUTE_IMMEDIATE	Executes non-query without bind variables

Procedure/Function	Description
FETCH_ROW	Fetches rows from query
GET_VALUE	Retrieves column value from SELECT statement or retrieves OUT bind parameters

This section contains these topics:

- [Executing Non-Queries](#)
- [Executing Queries](#)

### Executing Non-Queries

Non-queries include the following statements and types of statements:

- INSERT
- UPDATE
- DELETE
- DDL

To execute non-query statements, use the `EXECUTE_IMMEDIATE` function. For example, to execute a DDL statement at a non-Oracle system that you can access using the database link `SalesDB`, execute:

```
DECLARE
    num_rows INTEGER;

BEGIN
    num_rows := DBMS_HS_PASSTHROUGH.EXECUTE_IMMEDIATE@SalesDB
        ('CREATE TABLE DEPT (n SMALLINT, loc CHARACTER(10))');
END;
```

The variable `num_rows` is assigned the number of rows affected by the execution. For DDL statements, zero is returned. Note that you cannot execute a query with `EXECUTE_IMMEDIATE` and you cannot use bind variables.

**Using Bind Variables: Overview** Bind variables allow you to use the same SQL statement multiple times with different values, reducing the number of times a SQL statement needs to be parsed. For example, when you need to insert four rows in a particular table, you can parse the SQL statement once and bind and execute the SQL statement for each row. One SQL statement can have zero or more bind variables.


To execute pass-through SQL statements with bind variables, you must:

1. Open a cursor.
2. Parse the SQL statement at the non-Oracle system.
3. Bind the variables.
4. Execute the SQL statement at the non-Oracle system.
5. Close the cursor.

[Figure 3-1](#) shows the flow diagram for executing non-queries with bind variables.



**Figure 3–1** Flow Diagram for Non-Query Pass-Through SQL



**Using IN Bind Variables** The syntax of the non-Oracle system determines how a statement specifies a bind variable. For example, in Oracle you define bind variables with a preceding colon, as in:

```
UPDATE EMP  
SET SAL=SAL*1.1  
WHERE ENAME=:ename
```

In this statement, `ename` is the bind variable. In other non-Oracle systems you may need to specify bind variables with a question mark, as in:

```
UPDATE EMP
SET SAL=SAL*1.1
WHERE ENAME= ?
```

In the bind variable step, you must positionally associate host program variables (in this case, PL/SQL) with each of these bind variables.

For example, to execute the above statement, you can use the following PL/SQL program:

```
DECLARE
  c INTEGER;
  nr INTEGER;
BEGIN
  c := DBMS_HS_PASSTHROUGH.OPEN_CURSOR@SalesDB(
    DBMS_HS_PASSTHROUGH.PARSE@SalesDB(c,
      'UPDATE EMP SET SAL=SAL*1.1 WHERE ENAME=?');
  DBMS_HS_PASSTHROUGH.BIND_VARIABLE(c,1,'JONES');
  nr:=DBMS_HS_PASSTHROUGH.EXECUTE_NON_QUERY@SalesDB(c);
  DBMS_OUTPUT.PUT_LINE(nr||' rows updated');
  DBMS_HS_PASSTHROUGH.CLOSE_CURSOR@salesDB(c);
END;
```

**Using OUT Bind Variables** In some cases, the non-Oracle system can also support OUT bind variables. With OUT bind variables, the value of the bind variable is not known until *after* the execution of the SQL statement.

Although OUT bind variables are populated after the SQL statement is executed, the non-Oracle system must know that the particular bind variable is an OUT bind variable *before* the SQL statement is executed. You must use the `BIND_OUT_VARIABLE` procedure to specify that the bind variable is an OUT bind variable.

After the SQL statement is executed, you can retrieve the value of the OUT bind variable using the `GET_VALUE` procedure.


**Using IN OUT Bind Variables** A bind variable can be both an IN and an OUT variable. This means that the value of the bind variable must be known before the SQL statement is executed but can be changed after the SQL statement is executed.

For IN OUT bind variables, you must use the `BIND_INOUT_VARIABLE` procedure to provide a value *before* the SQL statement is executed. *After* the SQL statement is executed, you must use the `GET_VALUE` procedure to retrieve the new value of the bind variable.

## Executing Queries

The difference between queries and non-queries is that queries retrieve a result set from a `SELECT` statement. The result set is retrieved by iterating over a cursor.

[Figure 3-2](#) illustrates the steps in a pass-through SQL query. After the system parses the `SELECT` statement, each row of the result set can be fetched with the `FETCH_ROW` procedure. After the row is fetched, use the `GET_VALUE` procedure to retrieve the select list items into program variables. After all rows are fetched you can close the cursor.

**Figure 3–2 Pass-through SQL for Queries**

You do not have to fetch all the rows. You can close the cursor at any time after opening the cursor, for example, after fetching a few rows.

---

---

**Note:** Although you are fetching one row at a time, Heterogeneous Services optimizes the round trips between the Oracle9i server and the non-Oracle system by buffering multiple rows and fetching from the non-Oracle data system in one round trip.

---

---

The next example executes a query:

```
DECLARE
    val VARCHAR2(100);
    c    INTEGER;
    nr   INTEGER;
BEGIN
    c := DBMS_HS_PASSTHROUGH.OPEN_CURSOR@SalesDB;
    DBMS_HS_PASSTHROUGH.PARSE@SalesDB(c,
        'select ename
         from   emp
         where  deptno=10');
    LOOP
        nr := DBMS_HS_PASSTHROUGH.FETCH_ROW@SalesDB(c);
        EXIT WHEN nr = 0;
        DBMS_HS_PASSTHROUGH.GET_VALUE@SalesDB(c, 1, val);
        DBMS_OUTPUT.PUT_LINE(val);
    END LOOP;
    DBMS_HS_PASSTHROUGH.CLOSE_CURSOR@SalesDB(c);
END;
```

After parsing the `SELECT` statement, the rows are fetched and printed in a loop until the function `FETCH_ROW` returns the value 0.

**Note Also:** For more information on Passthrough SQL, please see [Appendix C, "DBMS\\_HS\\_PASSTHROUGH for Pass-Through SQL"](#)

## Result Set Support

### Introduction

Various relational databases allow stored procedures to return result sets. In other words, stored procedures will be able to return one or more sets of rows. This is a relatively new feature for any database.

Traditionally, database stored procedures worked exactly like procedures in any high-level programming language. They had a fixed number of arguments which could be of types `in`, `out`, or `in-out`. If a procedure had `n` arguments, it could return at most `n` values as results. However, suppose that somebody wanted a stored procedure to execute a query such as `SELECT * FROM emp` and return the results. The `emp` table might have a fixed number of columns but there is no way of telling, at procedure creation time, the number of rows it has. Because of this, no traditional stored procedure can be created that can return the results of a such a query. As a result, several relational database vendors added the capability of returning results sets from stored procedures, but each kind of relational database returns result sets from stored procedures in a different way.

Oracle has a data type called a ref cursor. Like every other Oracle data type, a stored procedure can take this data type as an `in` or `out` argument. In Oracle, a stored procedure can return a result set in the following way. To return a result set, a stored procedure must have an output argument of type `ref cursor`. It then opens a cursor for a SQL statement and places a handle to that cursor in that output parameter. The caller can then fetch from the ref cursor the same way as from any other cursor.

Oracle can do a lot more than simply return result sets. Ref cursors can be passed as input arguments to PL/SQL routines to be passed back and forth between client programs and PL/SQL routines or between several PL/SQL routines. Until recently, ref cursors in Oracle did not work in a distributed environment. This meant that you could pass ref cursor values between PL/SQL routines in the same database or between a client program and a PL/SQL routine, but they could not be passed from one database to another. As of Oracle9i, that restriction has been removed in the case of Heterogeneous Services.

### Result Set Support In Non-Oracle Systems:

Several non-Oracle systems allow stored procedures to return result sets but do so in completely different ways. No other relational database management system (RDBMS) has anything like the Oracle ref cursor data type. Result sets are supported to some extent in DB2, Sybase, Microsoft SQL Server, and Informix. Result set support in these databases is based on one of the following two models.

## Model 1

When creating a stored procedure, the user can explicitly specify the maximum number of result sets that can be returned by that stored procedure. While executing, the stored procedure can open anywhere from zero to its pre-specified maximum number of result sets. After the execution of the stored procedure, a client program can obtain handles to these result sets by using either an embedded SQL directive or calling a client library function. After that the client program can fetch from the result in the same way as from a regular cursor.

## Model 2

In this model, there is no pre-specified limit to the number of result sets that can be returned by a stored procedure. Both Model 1 and Oracle have a limit. For Oracle the number of result sets returned by a stored procedure can be at most the number of ref cursor `out` arguments; for Model 1, the upper limit is specified using a directive in the stored procedure language. Another way that Model 2 differs from Oracle and Model 1 is that they do not return a handle to the result sets but instead place the entire result set on the wire when returning from a stored procedure. For Oracle, the handle is the ref cursor `out` argument; for Model 1, it is obtained separately after the execution of the stored procedure. For both Oracle and Model 1, once the handle is obtained, data from the result set is obtained by doing a fetch on the handle; we have a bunch of cursors open and can fetch in any order. In the case of Model 2, however, all the data is already on the wire, with the result sets coming in the order determined by the stored procedure and the output arguments of the procedures coming at the end. So the whole of the first result set must be fetched, then the whole of the second one, until all of the results have been fetched. Finally, the stored procedure `out` arguments must be fetched.

## Heterogeneous Services Support for Result Sets

As can be seen in the preceding sections, result set support exists among non-Oracle databases in a variety of forms. All of these have to be mapped onto the Oracle ref cursor model. Due to the considerable differences in behavior among the various non-Oracle systems, Heterogeneous Services result set support will have to behave in one of two different ways depending on the non-Oracle system it is connected to.

Please note the following about Heterogeneous Services result set support:

- Result set support is present in 9i Heterogeneous Services generic code but in order for the feature to work in a gateway, the driver has to implement it as well. Not all drivers have implemented result set support and the customer must check in his gateway-specific documentation to determine whether it is supported in that gateway.

- Heterogeneous Services will support `ref cursor out` arguments from stored procedures. `In` and `in-out` arguments will not be supported.
- The `ref cursor out` arguments will all be anonymous ref cursors. No typed ref cursors are returned by Heterogeneous Services.

### Cursor mode

Oracle generally behaves such that each result set returned by the non-Oracle system stored procedure is mapped by the driver to an `out` argument of type `ref cursor`. The client program sees a stored procedure with several `out` arguments of type `ref cursor`. After executing the stored procedure, the client program can fetch from the `ref cursor` in exactly the same way as it would from a `ref cursor` returned by an Oracle stored procedure. When connecting to the gateway as described in Model 1, Heterogeneous Services will be in cursor mode.

### Sequential Mode

In Oracle, there is a pre-specified maximum number of result sets that a particular stored procedure can return. The number of result sets returned is at most the number of `ref cursor out` arguments for the stored procedure. It can, of course, return fewer result sets, but it can never return more.

For the system described in Model 2, there is no pre-specified maximum of result sets that can be returned. In the case of Model 1, we know the maximum number of result sets that a procedure can return, and the driver can return to Heterogeneous Services a description of a stored procedure with that many `ref cursor out` arguments. If, on execution of the stored procedure, fewer result sets than the maximum are returned, then the other `ref cursor out` arguments will be set to `NULL`.

Another problem for Model 2 database servers is that result sets have to be retrieved in the order in which they were placed on the wire by the database. This prevents Heterogeneous Services from running in cursor mode when connecting to these databases. To access result sets returned by these stored procedures, you must operate Heterogeneous Services in sequential mode.

In sequential mode, the procedure description returned by the driver contains the following:

- All the input arguments of the remote stored procedure
- None of the output arguments
- One `out` argument of type `ref cursor` (corresponding to the first result set returned by the stored procedure)



The client fetches from this ref cursor and then calls the virtual package function `dbms_`  
`hs_result_set.get_next_result_set` to get the ref cursor corresponding to the  
next result set. This function call is repeated until all result sets have been fetched. The last  
result set returned will actually be the `out` arguments of the remote stored procedure.

The major limitations of sequential mode are as follows:

- Result sets returned by a remote stored procedure have to be retrieved in the order in which they were placed on the wire
- On execution of a stored procedure, all result sets returned by a previously executed stored procedure will be closed (regardless of whether the data has been completely fetched or not).

## Code Examples:

All examples in this section use the following non-Oracle system stored procedure.

---



---

**Note:** For purposes of illustration, the following examples are presented as if they were Oracle PL/SQL stored procedures. However, you can create equivalent stored procedures for the DB2, Microsoft SQL Server, and Sybase.

---



---

```
create or replace package rcpackage is
  type rctype is ref cursor;
end rcpackage;
/
```

```
create or replace procedure refcurproc
  (arg1 in varchar2, arg2 out varchar2,
   rc1 out rcpackage.rctype,
   rc2 out rcpackage.rctype) is
begin
  arg2 := arg1;
  open rc1 for select * from emp;
  open rc2 for select * from dept;
end;
/
```

This stored procedure assigns the input parameter `arg1` to the output parameter `arg2`, opens the query `SELECT * FROM emp` in ref cursor `rc1`, and opens the query `SELECT * FROM dept` in ref cursor `rc2`.

## OCI program fetching from result sets in cursor mode

The following example shows OCI program fetching from result sets in cursor mode.

```
OCIEnv *ENVH;
OCISvcCtx *SVCH;
OCIStmt *STMH;
OCIError *ERRH;
OCIBind *BNDH[4];
OraText arg1[20];
OraText arg2[20];
OCIResult *arg3, *arg4;
OCIStmt *rstmt1, *rstmt2;
ub2 rcode[4];
ub2 rlens[4];
sb2 inds[4];
OraText *stmt = (OraText *) "begin refcurproc@link(:1,:2,:3,:4); end;";

/* Handle Initialization code skipped */

/* Prepare procedure call statement */

OCIStmtPrepare(STMH, ERRH, stmt, strlen(stmt), OCI_NTV_SYNTAX,
OCI_DEFAULT);

/* Bind procedure arguments */

inds[0] = 0;
strcpy((char *) arg1, "Hello World");
rlens[0] = strlen(arg1);
OCIBindByPos(STMH, &BNDH[0], ERRH, 1, (dvoid *) arg1, 20, SQLT_CHR,
(dvoid *) &(inds[0]), &(rlens[0]), &(rcode[0]),
0, (ub4 *) 0, OCI_DEFAULT);

inds[1] = 0;
rlens[1] = 0;
OCIBindByPos(STMH, &BNDH[1], ERRH, 2, (dvoid *) arg2, 20, SQLT_CHR,
(dvoid *) &(inds[1]), &(rlens[1]), &(rcode[1]),
0, (ub4 *) 0, OCI_DEFAULT);

inds[2] = 0;
rlens[2] = 0;
OCIDescriptorAlloc(ENVH, (dvoid **) &arg3, OCI_DTYPE_RSET, 0,
(dvoid **) 0);
OCIBindByPos(STMH, &BNDH[2], ERRH, 3, (dvoid *) arg3, 0, SQLT_RSET,
(dvoid *) &(inds[2]), &(rlens[2]), &(rcode[2]),
```

```

        0, (ub4 *) 0, OCI_DEFAULT);

inds[3] = 0;
rlens[3] = 0;
OCIDescriptorAlloc(ENVH, (dvoid **) &arg4, OCI_DTYPE_RSET, 0,
                  (dvoid **) 0);
OCIBindByPos(SIMH, &BNDH[3], ERRH, 4, (dvoid *) arg4, 0, SQLT_RSET,
            (dvoid *) &(inds[3]), &(rlens[3]), &(rcode[3]),
            0, (ub4 *) 0, OCI_DEFAULT);

/* Execute procedure */

OCISstmtExecute(SVCH, SIMH, ERRH, 1, 0, (CONST OCISnapshot *) 0,
               (OCISnapshot *) 0, OCI_DEFAULT);

/* Convert result set descriptors to statement handles */

OCIResultSetToStmnt(arg3, ERRH);
OCIResultSetToStmnt(arg4, ERRH);
rstmt1 = (OCISstmt *) arg3;
rstmt2 = (OCISstmt *) arg4;

/* After this the user can fetch from rstmt1 and rstmt2 */

```

## OCI program fetching from result sets in sequential mode

The following example shows OCI program fetching from result sets in sequential mode.

```

OCIEnv *ENVH;
OCISvcCtx *SVCH;
OCISstmt *SIMH;
OCIError *ERRH;
OCIBind *BNDH[2];
OraText arg1[20];
OCIResult *rset;
OCISstmt *rstmt;
ub2 rcode[2];
ub2 rlens[2];
sb2 inds[2];
OraText *stmt = (OraText *) "begin refcurproc@link(:1,:2); end;";
OraText *n_rs_stm = (OraText *)
    "begin :ret := DBMS_HS_RESULT_SET.GET_NEXT_RESULT_SET@link; end;";

/* Prepare procedure call statement */

```

```
/* Handle Initialization code skipped */

OCIStmtPrepare(STMH, ERRH, stmt, strlen(stmt), OCI_NTV_SYNTAX,
              OCI_DEFAULT);

/* Bind procedure arguments */

inds[0] = 0;
strcpy((char *) arg1, "Hello World");
rlens[0] = strlen(arg1);
OCIBindByPos(STMH, &BNDH[0], ERRH, 1, (dvoid *) arg1, 20, SQLT_CHR,
             (dvoid *) &(inds[0]), &(rlens[0]), &(rcode[0]),
             0, (ub4 *) 0, OCI_DEFAULT);

inds[1] = 0;
rlens[1] = 0;
OCIDescriptorAlloc(ENVH, (dvoid **) &rset, OCI_DTYPE_RSET, 0,
                  (dvoid **) 0);
OCIBindByPos(STMH, &BNDH[1], ERRH, 2, (dvoid *) rset, 0, SQLT_RSET,
             (dvoid *) &(inds[1]), &(rlens[1]), &(rcode[1]),
             0, (ub4 *) 0, OCI_DEFAULT);

/* Execute procedure */

OCIStmtExecute(SVCH, STMH, ERRH, 1, 0, (CONST OCISnapshot *) 0,
              (OCISnapshot *) 0, OCI_DEFAULT);

/* Convert result set to statement handle */

OCIResultSetToStmnt(rset, ERRH);
rstmt = (OCIStmt *) rset;

/* After this the user can fetch from rstmt */

/* Issue get_next_result_set call to get handle to next_result set */

/* Prepare Get next result set procedure call */

OCIStmtPrepare(STMH, ERRH, n_rs_stm, strlen(n_rs_stm), OCI_NTV_SYNTAX,
              OCI_DEFAULT);

/* Bind return value */

OCIBindByPos(STMH, &BNDH[1], ERRH, 1, (dvoid *) rset, 0, SQLT_RSET,
```

```

        (dvoid *) &(inds[1]), &(rlens[1]), &(rcode[1]),
        0, (ub4 *) 0, OCI_DEFAULT);

/* Execute statement to get next result set*/

OCIStmtExecute(SVCH, STMH, ERRH, 1, 0, (CONST OCISnapshot *) 0,
              (OCISnapshot *) 0, OCI_DEFAULT);

/* Convert next result set to statement handle */

OCIResultSetToStmt(rset, ERRH);
rstmt = (OCIStmt *) rset;

/* Now rstmt will point to the second result set returned by the
   remote stored procedure */

/* Repeat execution of get_next_result_set to get the output
   arguments */

```

## PL/SQL program fetching from result sets in cursor mode

Assume that the table `loc_emp` is a local table exactly like the remote `emp` table. The same assumption applies for `loc_dept`.

```

declare
    rc1    rcpackage.rctype;
    rec1   loc_emp%rowtype;
    rc2    rcpackage.rctype;
    rec2   loc_dept%rowtype;
    arg2   varchar2(20);

begin

    -- Execute procedure

    refcurproc@link('Hello World', arg2, rc1, rc2);

    -- Fetch 20 rows from the remote emp table and insert them
    -- into loc_emp

    for i in 1 .. 20 loop
        fetch rc1 into rec1;
        insert into loc_emp (rec1.empno, rec1.ename, rec1.job,
                           rec1.mgr, rec1.hiredate, rec1.sal,

```

```
                rec1.comm, rec1.deptno);
end loop;

-- Close the ref cursor

close rc1;

-- Fetch 5 rows from the remote dept table and insert them
-- into loc_dept

for i in 1 .. 5 loop
    fetch rc2 into rec2;
    insert into loc_dept values (rec2.deptno, rec2.dname, rec2.loc);
end loop;

-- Close the ref cursor

close rc2;

end;
```

PL/SQL program fetching from result sets in sequential mode

loc\_emp and loc\_dept are same as above. outarguments is a table with columns corresponding to the out arguments of the remote stored procedure

```
declare
    rc1   rcpackage.rctype;
    rec1  loc_emp%rowtype;
    rc2   rcpackage.rctype;
    rec2  loc_dept%rowtype;
    rc3   rcpackage.rctype;
    rec3  outargs%rowtype;

begin

    -- Execute procedure

    refcurproc@link('Hello World', rc1);

    -- Fetch 20 rows from the remote emp table and insert them
    -- into loc_emp

    for i in 1 .. 20 loop
        fetch rc1 into rec1;
```

```
        insert into loc_emp (rec1.empno, rec1.ename, rec1.job,
                           rec1.mgr, rec1.hiredate, rec1.sal,
                           rec1.comm, rec1.deptno);
    end loop;

    -- Close ref cursor

    close rc1;

    -- Get the next result set returned by the stored procedure

    rc2 := dbms_hs_result_set.get_next_result_set@link;

    -- Fetch 5 rows from the remote dept table and insert them
    -- into loc_dept

    for i in 1 .. 5 loop
        fetch rc2 into rec2;
        insert into loc_dept values (rec2.deptno, rec2.dname, rec2.loc);
    end loop;

    --Close ref cursor

    close rc2;

    -- Get the output arguments from the remote stored procedure
    -- Since we are in sequential mode, they will be returned in the
    -- form of a result set

    rc3 := dbms_hs_result_set.get_next_result_set@link;

    --Fetch them and insert them into the outarguments table

    fetch rc3 into rec3;
    insert into outarguments (rec3.col);

    --Close ref cursor

    close rc3;

end;
```

## Data Dictionary Translations

Most database systems have some form of data dictionary. A data dictionary is a collection of information about the database objects that have been created by various users of the system. For a relational database, a data dictionary is a set of tables and views which contain information about the data in the database. This information includes information on the users who are using the system and on the objects that they have created (such as tables, views, triggers and so forth). For the most part, all data dictionaries (regardless of the database system) contain the same information but each database system organizes the information in a different way.

For example, the Oracle data dictionary view `ALL_CATALOG` gives a list of tables, views, and sequences in the database. It has three columns: the first is called `OWNER` and is the name of the owner of the object, the second is called `TABLE_NAME` and is the name of the object, and the third is called `TABLE_TYPE` and is the type. This field has value `TABLE`, `VIEW`, `SEQUENCE` and so forth depending on the object type. However, in Sybase, the same information is stored two tables called `sysusers` and `sysobjects` whose column names are quite different than those of Oracle `ALL_CATALOG` table. Additionally, in Oracle, the table type is a string with value `TABLE`, `VIEW` and so forth but in Sybase it is a letter. For example, in Sybase, `U` means user table, `S` means system table, `V` means view, and so forth.

If the client program wanted information from the table `ALL_CATALOG` at Sybase then all it would have to do is to send a query referencing `ALL_CATALOG@database link` to a gateway and Heterogeneous Services will translate this query to the appropriate one on `sysobjects` and send the translated query to Sybase.

```
select SU."name" OWNER, SO."name" TABLE_NAME,
       decode(SO."type", 'U ', 'TABLE', 'S ', 'TABLE', 'V ', 'VIEW')
TABLE_TYPE
from "dbo"."sysusers"@link SU, "dbo"."sysobjects"@link SO
where SU."uid" = SO."uid" and
      (SO."type" = 'V' or SO."type" = 'S' or SO."type" = 'U')>
```

To relay such a translation of a query on an Oracle data dictionary table to the equivalent one on the non-Oracle system data dictionary table, Heterogeneous Services needs data dictionary translations for that non-Oracle system. A data dictionary translation is a view definition (essentially a select statement) over one or more non-Oracle system data dictionary tables such that the view looks exactly like the Oracle data dictionary table, with the same column names and the same information formatting. A data dictionary translation need not be as simple as the one above. Often the information needed is not found in one or two tables but is scattered over many tables and the data dictionary translation is a complex join over those tables.



In some cases, an Oracle data dictionary table does not have a translation because the information needed does not exist at the non-Oracle system. In such cases, the gateway can decide not to upload a translation at all or can resort to an alternative approach called **mimicking**. If the gateway wants to mimic a data dictionary table then it will let Heterogeneous Services know and Heterogeneous Services will obtain the description of the data dictionary table by querying the local database but when asked to fetch data, it will report that no rows were selected.

## Examples

The examples given below show the output of some data dictionary queries sent to Informix, and they compare the results with those produced when querying the same view on Oracle.

---



---

**Note:** The following examples use Informix as the non-Oracle system.

---



---

### Example 1: Check current session's user name on Oracle and on Informix.

To check the current session's user name on Oracle and on Informix, enter the following:

```
SQL select a.username, b.username from user_users a, user_users@remote_db b;
```

USERNAME	USERNAME
-----	-----
THSU	thsu

---



---

**Note:** Oracle maintains usernames in uppercase, Informix maintains them in lowercase.

---



---

### Example 2: Check current session's user ID on Oracle and on Informix.

To check the current session's user ID on Oracle and on Informix, enter the following:

```
SQL select a.user_id, b.user_id from user_users a, user_users@remote_db b;
```

USER_ID	USER_ID
-----	-----
25	0

---

---

**Note:** The Informix user ID is defaulted to zero because Informix does not maintain numeric USER\_ID values. This also illustrates the need to use caution when accessing other information on the Oracle server. Even if the connected non-Oracle system returns a value for its equivalent of USER\_ID, this is not a USER\_ID that is meaningful to the Oracle server since it applies only to the non-Oracle system. It would not be meaningful to do other Oracle data dictionary queries using the non-Oracle USER\_ID as a key.

---

---

**Example 3: Check constraints defined on a non-Oracle system for tables owned by an arbitrary user.**

To check constraints defined on a non-Oracle system for tables owned by an arbitrary user, enter the following:

```
SQL select constraint_name, table_name from all_constraints@remote_db
      2 where owner = 'thsu';
```

CONSTRAINT_NAME	TABLE_NAME
u19942_5270	thsmv_order_line
u24612_7116	thsmv_customer
u24613_7117	thsmv_orders

---

---

**Note:** Informix uses a different form of constraint names than Oracle, and its data dictionary maintains the table names in lowercase instead of uppercase.

---

---

**See Also:** For more information on data dictionary translations, please see [Appendix D, "Data Dictionary Translation Support"](#)

## Date Time

Oracle has five date time data types:

- Timestamp
- Timestamp with timezone
- Timestamp with local timezone
- Interval year to month
- Interval day to second

Heterogeneous Services generic code supports Oracle datetime data types in SQL and stored procedures. Oracle does not support these data types in data dictionary translations or queries involving data dictionary translations.

Even though Heterogeneous Services generic code supports this, support for a particular gateway depends on whether or not the driver for that non-Oracle system has implemented datetime support. Support even when the driver implements it may be partial because of the limitations of the non-Oracle system. Users should consult the documentation for their particular gateway on this issue.

The user must set the timestamp formats of the non-Oracle system in the gateway initialization file. The parameters to set are `HS_NLS_TIMESTAMP_FORMAT` and `HS_NLS_TIMESTAMP_TZ_FORMAT`. The user should also set the local time zone for the non-Oracle system in the initialization file. Parameter to set is `HS_TIME_ZONE`.

**See Also:** Oracle9i SQL Reference for information on datetime data types

## Two Phase Commit Protocol

Heterogeneous Services provides the infrastructure for the implementation of the two-phase commit mechanism. The extent to which this is supported depends on the gateway, and the remote system. Please refer to individual gateway manuals for more information.

**See Also:** For more information on two-phase commit protocol, see "Managing Distributed Transactions" in the *Oracle9i Administrator's Guide*.

## Piecewise Long

Earlier versions of gateways had limited support for the LONG data type. LONG is an Oracle data type that can be used to store up to 2 gigabytes (GB) of character/raw data (LONG RAW). These earlier versions restricted the amount of LONG data to 4 MB. This was because they would treat LONG data as a single piece. This led to restrictions of memory and network bandwidth on the size of the data that could be handled. Current gateways have extended the functionality to support the full 2 GB of heterogeneous LONG data. They handle the data piecewise between the agent and the Oracle server, thereby doing away with the large memory and network bandwidth requirements.

There is a new Heterogeneous Services initialization parameter, `HS_LONG_PIECE_TRANSFER_SIZE`, that can be used to set the size of the transferred pieces. For example, let us consider fetching 2 GB of LONG data from a heterogeneous source. A smaller piece size means less memory requirement, but more round trips to fetch all the data. A larger piece size means fewer round trips, but more of a memory requirement to store the intermediate pieces internally. Thus, the initialization parameter can be used to tune a system for the best performance, that is, for the best trade-off between round-trips and memory requirements. If the initialization parameter is not set, the system defaults to a piece size of 64 KB.

---

---

**Note:** This feature is not to be confused with piecewise operations on LONG data on the client side. Piecewise fetch and insert operations on the client side did work with the earlier versions of the gateways, and continue to do so. The only difference on the client side is that, where earlier versions of the gateways were able to fetch only up to 4 megabytes (MB) of LONG data, now they can fetch the entire 2 GB of LONG data. This is a significant improvement, considering that 4 MB is only 0.2% of the data type's full capacity.

---

---

## SQL\*Plus Describe Command

Until Oracle9i, you could not describe non-Oracle system objects using the SQL\*Plus DESCRIBE command. As of Oracle9i, functionality to do this has been added to Heterogeneous Services. There are still some limitations. For instance, using Heterogeneous links, you still cannot describe packages, sequences, synonyms, or types.

The SQL\*Plus DESCRIBE command is implemented using the OCIDescribeAny call, which was likewise unavailable before Oracle9i. The OCIDescribeAny call can also describe databases and schemas, which you cannot do through the SQL\*Plus DESCRIBE command. With Heterogeneous Services, you can do both.

In order to implement this functionality some additional driver logic is needed; not all drivers may have implemented it. Please consult individual gateway documentation to see if this feature is supported in that gateway.

## Constraints on SQL in a Distributed Environment

This section explains some of the constraints that exist on SQL in a distributed environment. These constraints apply to distributed environments that involve access to non-Oracle systems or remote Oracle databases.

This section contains the following topics:

- [Resolving Remote and Heterogeneous References](#)
- [Resolving Important Restrictions](#)
- [Updates, Inserts and Deletes](#)

## Resolving Remote and Heterogeneous References

---

---

**Note:** Many of the rules for Heterogeneous access also apply to remote references. For more information, please see the distributed database section of the *Oracle9i Database Administrator's Guide*.

---

---

A statement can, with restrictions, be executed on any database node referenced in the statement or the local node. If all objects referenced are resolved to a single, referenced node, then Oracle will attempt to execute a query at that node. You can force execution at a referenced node by using the /\*+ REMOTE\_MAPPED \*/ or /\*+ DRIVING\_SITE \*/ hints. If a statement is forwarded to a different node than the node it was issued at, then the statement is said to be **remote mapped**.

The ways in which statements can, must, and cannot be remote mapped are subject to specific rules or restrictions. If these rules are not all followed, then an error will occur. As long as the statements issued are consistent with all these rules, the order in which the rules are applied does not matter.

Different constraints exist when you are using SQL for remote mapping in a distributed environment. This distributed environment can include remote Oracle databases as well as databases that involve Oracle Transparent Gateways or Generic Connectivity connections between Oracle and non-Oracle systems.

## Resolving Important Restrictions

The following section lists some of the different constraints that exist when you are using SQL for remote mapping in a distributed environment.

---

---

**Note:** In the examples that follow, *remote\_db* refers to a remote non-Oracle system while *remote\_oracle\_db* refers to a remote Oracle server.

---

---

### **Rule A: A data definition language statement cannot be remote mapped.**

In Oracle data definition language, the target object syntactically has no place for a remote reference. Data definition language statements that contain remote references are always executed locally. For Heterogeneous Services, this means it cannot create a database for the non-Oracle database directly using SQL.

However, there is an indirect way using passthrough SQL.

Consider the following example:

```
begin
  dbms_hs.passthroughsql.execute_immediate@remote_db
  (
    'create table x1 (c1 char, c2 number)'
  );
end;
```

### **Rule B: Insert, Update and Delete statements with a remote target table must be remote mapped.**

This rule is more restrictive for non-Oracle remote databases than for a remote Oracle database. This is because the remote system cannot fetch data from the originating Oracle database while executing DML statements targeting tables in a non-Oracle system.

For example, to insert all local employees from the local emp1 table to a remote Oracle emp2 table, use the following statement:

```
INSERT INTO emp2@remote_oracle_db SELECT * FROM emp1;
```

This statement is remote mapped to the remote database. The remote mapped statement sent to the remote database contains a remote reference back to the originating database for emp1. Such a remote link received by the remote database is called a callback link.

In general however, gateways callback links are not supported. When you try to insert into a non-Oracle system using a select statement referencing a local table, an error occurs.

For example, consider the following statement:

```
INSERT INTO emp2@remote_db SELECT * from emp1;
```

The statement returns the following error message:

```
ORA-02025: all tables in the SQL statement must be at the remote database
```

The work around is to write a PL/SQL block:

```
declare
cursor remote_insert is select * from emp2;
begin

    for rec in remote_insert loop
        insert into emp1@remote_db (empno, ename, deptno) values (
            rec.empno,
            rec.ename,
            rec.deptno
        );
    end loop;
end;
/
```

Another special case are session specific SQL functions such as USER, USERENV and SYSDATE. These functions may need to be executed at the originating site. A remote mapped statement containing these functions will contain a callback link. For a non-Oracle database where callbacks are not supported this could (by default) result in a restriction error.

For example, consider the following statement:

```
DELETE FROM emp1@remote_db WHERE hiredata > sysdate;
```

The statement returns the following error message:

ORA-02070: database *REMOTE\_DB* does not support special functions in this context

This often must be resolved by replacing special functions with a bind variable:

```
DELETE FROM empl@remote_db WHERE hiredata > :1
```

**Rule C: Object features like tables with nested table columns, ADT columns, Opaque columns or Ref Columns cannot be remote mapped.**

Currently, the above column types are not supported for heterogeneous access. Hence, this limitation is not directly encountered.

**Rule D: SQL statements containing operators and constructs that are not supported at the remote site cannot be remote mapped.**

Note that in our description of Rule B we already encountered special constructs such as callback links and special functions as examples of this.

If the statement is a *select* (or *dml* with the target table local) and none of the remaining rules would require the statement to be remote mapped the statement can still be executed by processing the query locally using the local SQL engine and the remote *select* operation.

The remote *select* operation is the operation to retrieve rows for remote table data as opposed to other operations like full table scan and index access which retrieve rows of local table data. The remote table scan has a SQL statement associated with the operation. A full table scan of table *empl* is issued as `SELECT * FROM empl` (with the *\** expanded to the full column list). Access for indexes is converted back to where clause predicates and also filters that can be supported are passed down to the *WHERE* clause of the remote row source.

You can check the SQL statement generated by the Oracle server by explaining the statement and querying the *OTHER* column of the explain plan table for each *REMOTE* operation.

**See Also:** [Using Index and Table Statistics](#) for more information on how to interpret explain plans with remote references.

For example consider the following statement:

```
SELECT COUNT(*) FROM empl@remote_db WHERE hiredate < sysdate;
```

The statement returns the following output:

```
COUNT(*)
-----
      14
```



1 row selected.

The remote table scan is:

```
SELECT hiredate FROM empl
```

Since the predicate converted to a filter cannot be generated back and passed down to the remote operation because `sysdate` is not supported by the `remote_db` or evaluation rules, `sysdate` must be executed locally.

---

---

**Note:** Because the remote table scan operation is only partially related to the original query, the number of row retrieved can be significantly larger than you would expect and can have a significant impact on performance.

---

---

**Rule E: SQL statement containing a table expression cannot be remote mapped.**

This limitation is not directly encountered since table expressions are not supported in the heterogeneous access module.

**Rule F: If a SQL statement selects a long, the statement must be mapped to the node where the table containing the long resides.**

For example, consider the following statement:

```
SELECT long1 FROM table_with_long@remote_db, dual;
```

The statement returns the following error message:

```
ORA-02025: all tables in the SQL statement must be at the remote database
```

This can be resolved by the following statement:

```
SELECT long1 FROM table_with_long@remote_db WHERE long_idx = 1;
```

**Rule G: The statement must be mapped to the node on which the table or tables with columns referenced in the FOR UPDATE OF clause resides when the SQL statement is of form "SELECT...FOR UPDATE OF..."**

When the SQL statement is of the form `SELECT . . . FOR UPDATE OF . . .`, the statement must be mapped to the node on which the table or tables with columns referenced in the `FOR UPDATE OF` clause resides.

For example, consider the following statement:

```
SELECT ename FROM emp1@remote_db WHERE hiredate < sysdate FOR UPDATE OF empno
```

The statement returns the following error message:

```
ORA-02070: database REMOTE_DB does not support special functions in this context
```

**Rule H: If the SQL statement contains a SEQUENCE or sequences, the statement must be mapped to the site where each sequence resides.**

This rule is not encountered for the heterogeneous access since remote non-Oracle sequences are not supported. The restriction for remote non-Oracle access is already present because of the callback link restriction.

**Rule I: If the statement contains a user defined operator or operators, the statement must be mapped to the node where each operator is defined.**

This rule is also already covered under the callback link restriction discussed in Rule B.

**Rule J: A statement containing duplicate bind variables cannot be remote mapped.**

The work around for this restriction is to use unique bind variables and bind by number.

## Updates, Inserts and Deletes

As discussed in the previous section, updates to remote non-Oracle objects through an Oracle server are restricted by the missing callback feature support present in the Oracle database. This restricts data manipulation language (DML) upon remote non-Oracle database objects to statements that reference all objects in that remote non-Oracle database or are literals or bind variables.

Because of this, no objects can be referenced from the originating Oracle server or other remote objects.

Also, as with any remote update, whether non-Oracle or a previous remote update, if a SQL update in an Oracle format is not supported, then an error is returned in the following format:

```
ORA-2070: database ... does not support ... in this context.
```

---

---

**Note:** These restrictions do not apply to DML with a local target object referencing non-Oracle or remote Oracle database objects.

---

---

You can perform DML to remote Oracle or non-Oracle target tables in an Oracle format that is not supported by using PL/SQL. Declare a cursor that selects the appropriate row and executes the update for each row selected. The row may need to be unique, identified by selecting a primary key, or, if not available, a rowid.

Consider the following example:

```
declare
  v_empno number;
  cursor remote_update is select empno from emp1@remote_db
                        where  ename = v_ename;
  cursor c1 is select ename from emp2 where comm IS NOT NULL;
begin
  for recl in c1 loop
    v_ename = recl.ename;
    for rec in remote_update loop
      update emp1@remote_db set comm = 100 where empno rec.empno;
    end loop;
  end loop;
end;
/
```

## Using Index and Table Statistics

Heterogeneous Services collects certain table and index statistics information on the respective non-Oracle system tables and passes this information back to the Oracle server. The Oracle cost based optimizer uses this information when building the query plan.

For example consider the following statement where you create a table in the Oracle database with 10 rows:

```
CREATE table_T1 (C1 number);
```

Analyze the table by issuing the following SQL statement:

```
ANALYZE table_T1 COMPUTE STATISTICS;
```

Now create a table in the non-Oracle system with 1000 rows:

```
CREATE TABLE remote_t1 (C1 number)
```

Issue the following SQL statement:

```
SELECT a.* FROM remote_t1@remote_db a, T1 b
       where a.C1 = b.C1
```

The Oracle optimizer issues the following SQL statement to the agent:

```
SELECT C1 FROM remote_t1
```

This fetches all the 1000 rows from the non-Oracle system and performs the join in the Oracle database.

Now, if we add a unique index on the column C1 in the table `remote_t1`, and issue the same SQL statement again, the agent receives the following SQL statement:

```
SELECT C1 FROM remote_t1 WHERE C1 = ?
```

for each value of C1 in the local `t1`.

---

---

**Note:** ('?') is the bind parameter marker. Also, join predicates containing bind variables generated by Oracle are only generated for nested loop join methods.

---

---

To verify the SQL execution plan, generate an explain plan for the SQL statement. Load `utlxplan` in the `admin` directory first.

At the command prompt, type:

```
Explain plan for SELECT a.* FROM remote_t1@remote_db a, T1 b
  where a.C1 = b.C1;
```

Then, run the `utlxpls` utility script by entering the following statement.

```
@utlxpls
```

The operation `remote` indicates that remote SQL is being referenced.

To find out what statement is sent, type the following statement at the command prompt:

```
select ID, OTHER from EXPLAIN_PLAN where OPERATION = 'REMOTE';
```

## Other Optimizations

There are several other optimizations that the cost based optimizer performs. The most important ones are remote sort elimination and remote joins.

## Remote Join Optimization

The following is an example of the remote join optimization capability of the Oracle database.

---



---

**Note:** The explain plan that uses tables from a non-Oracle system can differ from similar statements with local or remote Oracle table scans. This is because of the limitation on the statistics available to Oracle for non-Oracle tables. Most importantly, column selectivity is not available for non-unique indexes of non-Oracle tables. Because of the limitation of the statistics available, the following example is not necessarily what you encounter when doing remote joins for yourself and is intended for illustration only.

---



---

Consider the following example:

```
explain plan for
select e.ename, d.dname, f.ename, f.deptno from
dept d,
emp@remote_db e,
emp@remote_db f
  where e.mgr = f.empno
     and e.deptno = d.deptno
     and e.empno = f.empno;
```

@utlxpls

**Table 3–1 Explain Plan**

Operation	Name	Rows	Bytes	Cost	Pstar
SELECT STATEMENT		1	101	128	
HASH JOIN		2K	132K	19	
TABLE ACCESS FULL	DEPT	21	462	1	
REMOTE		2K	89K	16	

Issue the following statement:

```
SET longwidth 300
SELECT other FROM plan_table WHERE operation = 'REMOTE';
```

You get the following output:

```
SELECT
A1."ENAME",A1."MGR",A1."DEPTNO",A1."EMPNO",A2."ENAME",A2."DEPTNO",A2."EMPNO",A2.
"EMPNO" FROM "EMP" A1,"EMP" A2 WHERE A1."EMPNO"=A2."EMPNO" AND
A1."MGR"=A2."EMPNO"
```

## Optimizer Restrictions for non-Oracle Access

1. There are no column statistics for remote objects. This can result in poor execution plans. Verify the execution plan and use hints to improve the plan.
2. There is no optimizer hint to force a remote join. However, there is a remote query block optimization that can be used to rewrite the query slightly in order to get a remote join.

For instance, the earlier example can be rewritten to the form:

```
select v.ename, d.dname, d.deptno from
dept d,
  (select /*+ NO_MERGE */
   e.deptno deptno, e.ename ename emp@remote_db e, emp@remote_db f
   where e.mgr = f.empno
   and e.empno = f.empno;
 )
where v.deptno = d.deptno;
```

This guarantees a remote join because it has been isolated in a nested query with the NO\_MERGE hint.

---

## Using the Gateway

This chapter explains how to use Oracle Transparent Gateways.

This chapter contains the following sections:

- [Setting Up Access to Non-Oracle Systems](#)
- [Initialization Parameters](#)
- [Optimizing Data Transfers Using Bulk Fetch](#)
- [Registering Agents](#)
- [Oracle Database Server SQL Construct Processing](#)
- [Using Synonyms](#)
- [Copying Data from the Oracle Database Server to the Non-Oracle Database System](#)
- [Copying Data from the Non-Oracle Database System to the Oracle Database Server](#)
- [Heterogeneous Services Data Dictionary Views](#)
- [Using the Heterogeneous Services Dynamic Performance Views](#)

## Setting Up Access to Non-Oracle Systems

This section explains the generic steps to configure access to a non-Oracle system.

---

---

**Note:** The instructions for configuring your agent may differ slightly from the following instructions. Please see the *Installation and User's Guide* for your agent for more complete installation information.

---

---

The steps for setting up access to a non-Oracle system are:

- [Step 1: Install the Heterogeneous Services Data Dictionary](#)
- [Step 2: Set Up the Environment to Access Heterogeneous Services Agents](#)
- [Step 3: Create the Database Link to the Non-Oracle System](#)
- [Step 4: Test the Connection](#)

### Step 1: Install the Heterogeneous Services Data Dictionary

To install the data dictionary tables and views for Heterogeneous Services, you must run a script that creates all the Heterogeneous Services data dictionary tables, views, and packages. On most systems the script is called `catsh.sql` and resides in `$ORACLE_HOME/rdbms/admin`.

---

---

**Note:** The data dictionary tables, views, and packages may already be installed on your Oracle9i server. Check for the existence of Heterogeneous Services data dictionary views, for example, `SYS.HS_FDS_CLASS`.

---

---

### Step 2: Set Up the Environment to Access Heterogeneous Services Agents

To initiate a connection to the non-Oracle system, the Oracle9i server starts an agent process through the Oracle Net listener. For the Oracle9i server to be able to connect to the agent, you must:

1. Set up a Oracle Net service name for the agent that can be used by the Oracle9i server. The Oracle Net service name descriptor includes protocol-specific information needed to access the Oracle Net listener. The service name descriptor must include the `(HS=OK)` clause to make sure the connection uses Oracle9i Heterogeneous Services.



2. Set up the listener to listen for incoming request from the Oracle9i server and spawn Heterogeneous Services agents. Modify the `listener.ora` file so that the listener can start Heterogeneous Services agents, and then restart the listener.

### A Sample Entry for a Oracle Net Service Name

The following is a sample entry for the service name in the `tnsnames.ora` file:

```
Sybase_sales= (DESCRIPTION=
                (ADDRESS=(PROTOCOL=tcp)
                  (HOST=dlsun206)
                  (PORT=1521))
                (CONNECT_DATA = (SID=SalesDB))
                (Heterogeneous Services = OK))
```

The description of this service name is defined in `tnsnames.ora`, the Oracle Names server, or in third-party name servers using the Oracle naming adapter.

---



---

**Note:** Please see the *Installation and User's Guide* for your agent for more information about how to define the Oracle Net service name.

---



---

### A Sample Listener Entry

The following is a sample entry for the listener in `listener.ora`:

```
LISTENER =
  (ADDRESS_LIST =
    (ADDRESS= (PROTOCOL=tcp)
              (HOST = dlsun206)
              (PORT = 1521))
    )
  )
...
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC = (SID_NAME=SalesDB)
                (ORACLE_HOME=/home/oracle/megabase/8.1.3)
                (PROGRAM=tg4mb80))
    )
  )
```

The value associated with `PROGRAM` keyword defines the name of the agent executable. The agent executable must reside in the `$ORACLE_HOME/bin` directory. Typically, you use `SID_NAME` to define the initialization parameter file for the agent.

### Step 3: Create the Database Link to the Non-Oracle System

To create a database link to the non-Oracle system, use the `CREATE DATABASE LINK` statement. The **service name** that is used in the `USING` clause of the `CREATE DATABASE LINK` command is the Oracle Net service name.

For example, to create a database link to the `SALES` database on Sybase, enter:

```
CREATE DATABASE LINK sales
USING 'Sybase_sales';
```

### Step 4: Test the Connection

To test the connection to the non-Oracle system, use the database link in a SQL or PL/SQL statement. If the non-Oracle system is a SQL-based database, you can execute a `SELECT` statement from an existing table or view using the database link. For example, issue:

```
SELECT * FROM product@sales
WHERE product_name like '%pencil%';
```

When you try to access the non-Oracle system for the first time, the Heterogeneous Services agent uploads information into the Heterogeneous Services data dictionary. The uploaded information includes:

Type of Data	Explanation
Capabilities of the non-Oracle system	For example, the agent specifies whether it can perform a join, or a <code>GROUP BY</code> .
SQL translation information	The agent specifies how to translate Oracle functions and operators into functions and operators of the non-Oracle system.
Data dictionary translations	To make the data dictionary information of the non-Oracle system available just as if it were an Oracle data dictionary, the agent specifies how to translate Oracle data dictionary tables into tables and views of the non-Oracle system.

---

---

**Note:** Most agents upload information into the Oracle9i data dictionary automatically the first time they are accessed. Some agent vendors may provide scripts, however, that you must run on the Oracle9i server.

---

---

**See Also:** “Heterogeneous Services Data Dictionary Views” on page 4-22 and [Appendix D, "Data Dictionary Translation Support"](#).

## Initialization Parameters

As mentioned in “Configuring Heterogeneous Services” on page 2-5, the user can configure the gateway using initialization parameters. This is done by creating an initialization file and setting the desired parameters in this file

Heterogeneous Services parameters are distinct from Oracle database server initialization parameters. Heterogeneous Services initialization parameters are set in the Heterogeneous Services initialization file and not in the Oracle `init.ora` file. There is a Heterogeneous Services initialization file for each gateway instance. The name of the file is `initsid.ora`, where `sid` is the Oracle system identifier used for the gateway. In the case of generic connectivity, the file is located in the directory `$ORACLE_HOME/hs/admin` and in the case of transparent gateways it is located in the directory `$ORACLE_HOME/product_name/admin` where `product_name` is the name of the product. So, the Sybase gateway initialization file is located in the directory `$ORACLE_HOME/tg4sybs/admin`.

The syntax of the initialization file is as follows. The file contains a list of initialization parameter settings each of which should be on a separate line. The syntax to set an initialization parameter is:

```
[set] [private] parameter = value
```

The `set` and `private` keywords are optional. If the `set` keyword is present then the variable will also be set in the environment. If the `private` keyword is present, the parameter will not be uploaded to the server. In general, it is recommended that this keyword not be used - unless the initialization parameter value contains sensitive information (like a password) that should not be sent over the network from gateway to Oracle server.

Another initialization file can be included in an Heterogeneous Services initialization file by using the `ifile` directive. The syntax for this is

```
ifile = pathname for file to be included
```

In the initialization parameter syntax, all keywords (`SET`, `PRIVATE` and `IFILE`) are case insensitive. Initialization parameter names and values are case sensitive. Most initialization parameters names will be uppercase. When there are any exceptions to this rule, we will explicitly point them out.

Gateway initialization parameters can be divided into two groups. One is a set of generic initialization parameters that are common to all gateways and the other is a set of initialization parameters that are specific to individual gateways. The list of generic initialization parameters is given below. Please refer individual gateway documentation for the list of initialization parameters specific to that gateway.

---

---

**Note:** Most (but not all) gateway initialization parameter follow these conventions. For more information on initialization parameters for individual gateways, please see your gateway specific documentation.

---

---

- HS\_COMMIT\_POINT\_STRENGTH
- HS\_DB\_DOMAIN
- HS\_DB\_INTERNAL\_NAME
- HS\_DB\_NAME
- HS\_DESCRIBE\_CACHE\_HWM
- HS\_FDS\_CONNECT\_INFO
- HS\_FDS\_SHAREABLE\_NAME
- HS\_FDS\_TRACE\_LEVEL
- HS\_FDS\_TRACE\_FILE\_NAME
- HS\_LANGUAGE
- HS\_LONG\_PIECE\_TRANSFER\_SIZE
- HS-NLS\_DATE\_FORMAT
- HS-NLS\_DATE\_LANGUAGE
- HS-NLS\_NCHAR
- HS-NLS\_TIMESTAMP\_FORMAT
- HS-NLS\_TIMESTAMP\_TZ\_FORMAT
- HS\_OPEN\_CURSORS
- HS\_ROWID\_CACHE\_SIZE
- HS\_RPC\_FETCH\_REBLOCKING
- HS\_RPC\_FETCH\_SIZE
- HS\_TIME\_ZONE

Do not use the private keyword when setting any of these parameters. Doing that would cause the parameter not to be uploaded to the server and could cause errors

in SQL processing. None of these parameters need be set in the environment, so the set keyword need not be used either.

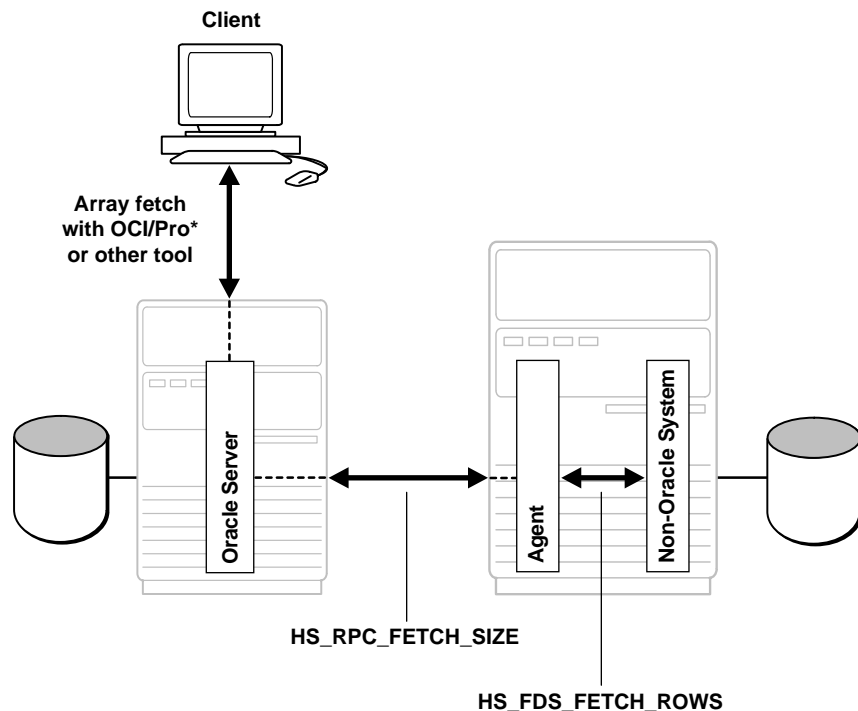
## Optimizing Data Transfers Using Bulk Fetch

When an application fetches data from a non-Oracle system using Heterogeneous Services, data is transferred:

- From the non-Oracle system to the agent process
- From the agent process to the Oracle database server
- From the Oracle database server to the application

Oracle allows you to optimize all three data transfers, as illustrated in [Figure 4-1](#).

*Figure 4-1 Optimizing data transfers*



This section contains the following topics:

- [Using OCI, an Oracle Precompiler, or Another Tool for Array Fetches](#)
- [Controlling the Array Fetch Between Oracle Database Server and Agent](#)
- [Controlling the Array Fetch Between Agent and Non-Oracle Server](#)
- [Controlling the Reblocking of Array Fetches](#)

## Using OCI, an Oracle Precompiler, or Another Tool for Array Fetches

You can optimize data transfers between your application and the Oracle9i server by using array fetches. See your application development tool documentation for information about array fetching and how to specify the amount of data to be sent per network round trip.

## Controlling the Array Fetch Between Oracle Database Server and Agent

When Oracle retrieves data from a non-Oracle system, the Heterogeneous Services initialization parameter `HS_RPC_FETCH_SIZE` defines the number of bytes sent per fetch between the agent and the Oracle9i server. The agent fetches data from the non-Oracle system until one of the following occurs:

- It has accumulated the specified number of bytes to send back to the Oracle database server.
- The last row of the result set is fetched from the non-Oracle system.

## Controlling the Array Fetch Between Agent and Non-Oracle Server

The initialization parameter `HS_FDS_FETCH_ROWS` determines the number of rows to be retrieved from a non-Oracle system. Note that the array fetch must be supported by the agent. See your agent-specific documentation to ensure that your agent supports array fetching.

## Controlling the Reblocking of Array Fetches

By default, an agent fetches data from the non-Oracle system until it has enough data retrieved to send back to the server. That is, it keeps going until the number of bytes fetched from the non-Oracle system is equal to or higher than the value of `HS_RPC_FETCH_SIZE`. In other words, the agent *reblocks* the data between the agent and the Oracle database server in sizes defined by the value of `HS_RPC_FETCH_SIZE`.

When the non-Oracle system supports array fetches, you can immediately send the data fetched from the non-Oracle system by the array fetch to the Oracle database server without waiting until the exact value of `HS_RPC_FETCH_SIZE` is reached. That is, you can stream the data from the non-Oracle system to the Oracle database server and disable reblocking by setting the value of initialization parameter `HS_RPC_FETCH_REBLOCKING` to `OFF`.

For example, assume that you set `HS_RPC_FETCH_SIZE` to 64 kilobytes (KB) and `HS_FDS_FETCH_ROWS` to 100 rows. Assume that each row is approximately 600 bytes in size, so that the 100 rows are approximately 60 KB. When `HS_RPC_FETCH_REBLOCKING` is set to `ON`, the agent starts fetching 100 rows from the non-Oracle system.

Because there is only 60 KB of data in the agent, the agent does not send the data back to the Oracle database server. Instead, the agent fetches the next 100 rows from the non-Oracle system. Because there is now 120 KB of data in the agent, the first 64 KB can be sent back to the Oracle database server.

Now there is 56 KB of data left in the agent. The agent fetches another 100 rows from the non-Oracle system before sending the next 64 KB of data to the Oracle database server. By setting the initialization parameter `HS_RPC_FETCH_REBLOCKING` to `OFF`, the first 100 rows are immediately sent back to the Oracle9i server.



## Registering Agents

**Registration** is an operation through which Oracle stores information about an agent in the data dictionary. Agents do not have to be registered. If an agent is not registered, Oracle stores information about the agent in memory instead of in the data dictionary; when a session involving an agent terminates, this information ceases to be available.

**Self-registration** is an operation in which a database administrator sets an initialization parameter that lets the agent automatically upload information into the data dictionary. In release 8.0 of the Oracle database server, an agent could determine whether to self-register. In release 9.0, self-registration occurs only when the `HS_AUTOREGISTER` initialization parameter is set to `TRUE` (default).

---



---

**Note:** `HS_AUTOREGISTER` is an Oracle initialization parameter that you set in the `init.ora` file; it is not a Heterogeneous Services initialization parameter that is set in the gateway initialization file.

---



---

This section contains the following topics:

- [Enabling Agent Self-Registration](#)
- [Disabling Agent Self-Registration](#)

### Enabling Agent Self-Registration

To ensure correct operation over heterogeneous database links, agent self-registration automates updates to Heterogeneous Services configuration data that describe agents on remote hosts. Agent self-registration is the default behavior. If you do not want to use the agent self-registration feature, then set the initialization parameter `HS_AUTOREGISTER` to `FALSE`.

Both the server and the agent rely on three types of information to configure and control operation of the Heterogeneous Services connection. These three sets of information are collectively called **HS configuration data**:

Heterogeneous Services Configuration Data	Description
Heterogeneous Services initialization parameters	Provide control over various connection-specific details of operation.

<b>Heterogeneous Services Configuration Data</b>	<b>Description</b>
Capability definitions	Identify details like SQL language features supported by the non-Oracle datasource.
Data dictionary translations	Map references to Oracle data dictionary tables and views into equivalents specific to the non-Oracle data source.

**See Also:** ["Specifying HS\\_AUTOREGISTER"](#) on page 4-14.

### Using Agent Self-Registration to Avoid Configuration Mismatches

HS configuration data is stored in the Oracle database server's data dictionary. Because the agent is possibly remote, and may therefore be administered separately, several circumstances can lead to configuration mismatches between servers and agents:

- An agent can be newly installed on a separate machine so that the server has no Heterogeneous Services data dictionary content to represent the agent's HS configuration data.
- A server can be newly installed and lack the necessary HS configuration data for existing agents and non-Oracle data stores.
- A non-Oracle instance can be upgraded from an older version to a newer version, requiring modification of the HS configuration data.
- An Heterogeneous Services agent at a remote site can be upgraded to a new version or patched, requiring modification of the HS configuration data.
- A database administrator (DBA) at the non-Oracle site can change the agent setup, possibly for tuning or testing purposes, in a manner which affects HS configuration data.

Agent self-registration permits successful operation of Heterogeneous Services in all these scenarios. Specifically, agent self-registration enhances inter operability between any Oracle database server and any Heterogeneous Services agent, provided that each is at least as recent as Version 8.0.3. The basic mechanism for this functionality is the ability to upload HS configuration data from agents to servers.

Self-registration provides automatic updating of HS configuration data residing in the Oracle database server data dictionary. This update ensures that the agent self-registration uploads need to be done only once, on the initial use of a previously unregistered agent. Instance information is uploaded on each connection, not stored in the server data dictionary.

### Understanding Agent Self-Registration

The Heterogeneous Services agent self-registration feature can:

- Identify the agent and the non-Oracle data store to the Oracle database server.
- Permit agents to define Heterogeneous Services initialization parameters for use both by the agent and connected Oracle9i servers.
- Upload capability definitions and data dictionary translations, if available, from an Heterogeneous Services agent during connection initialization.

---

---

**Note:** When both the server and the agent are release 8.1 or higher, the upload of class information occurs only when the class is undefined in the server data dictionary. Similarly, instance information is uploaded only if the instance is undefined in the server data dictionary.

---

---

The information required to accomplish the above is accessed in the server data dictionary by using these agent-supplied names:

- FDS\_CLASS
- FDS\_CLASS\_VERSION

**See Also:** ["Heterogeneous Services Data Dictionary Views"](#) on page 4-22 to learn how to use the Heterogeneous Services data dictionary views.

**FDS\_CLASS and FDS\_CLASS\_VERSION** FDS\_CLASS and FDS\_CLASS\_VERSION are defined by Oracle or by third-party vendors for each individual Heterogeneous Services agent and version. Oracle Heterogeneous Services concatenates these names to form FDS\_CLASS\_NAME, which is used as a primary key to access class information in the server data dictionary.

FDS\_CLASS should specify the type of non-Oracle data store to be accessed and FDS\_CLASS\_VERSION should specify a version number for both the non-Oracle

data store and the agent that connects to the it. Note that when any component of an agent changes, `FDS_CLASS_VERSION` must also change to uniquely identify the new release.

---

---

**Note:** This information is uploaded when you initialize each connection.

---

---

**FDS\_INST\_NAME Instance-specific information** can be stored in the server data dictionary. The instance name, `FDS_INST_NAME`, is configured by the DBA who administers the agent; how the DBA performs this configuration depends on the specific agent in use.

The Oracle database server uses `FDS_INST_NAME` to look up instance-specific configuration information in its data dictionary. Oracle uses the value as a primary key for columns of the same name in these views:

- `FDS_INST_INIT`
- `FDS_INST_CAPS`
- `FDS_INST_DD`

Server data dictionary accesses that use `FDS_INST_NAME` also use `FDS_CLASS_NAME` to uniquely identify configuration information rows. For example, if you port a database from class Sybase8.1.6 to class Sybase8.1.7, both databases can simultaneously operate with instance name `SCOTT` and use separate sets of configuration information.

Unlike class information, instance information is not automatically self-registered in the server data dictionary.

- If the server data dictionary contains instance information, it represents DBA-defined setup details which fully define the instance configuration. No instance information is uploaded from the agent to the server.
- If the server data dictionary contains no instance information, any instance information made available by a connected agent is uploaded to the server for use in that connection. The uploaded instance data is not stored in the server data dictionary.

### Specifying `HS_AUTOREGISTER`

The Oracle database server initialization parameter `HS_AUTOREGISTER` enables or disables automatic self-registration of Heterogeneous Services agents. Note that this

parameter is specified in the Oracle initialization parameter file, not the agent initialization file. For example, you can set the parameter as follows:

```
HS_AUTOREGISTER = TRUE
```

When set to `TRUE`, the agent uploads information describing a previously unknown agent class or a new agent version into the server's data dictionary.

Oracle recommends that you use the default value for this parameter (`TRUE`), which ensures that the server's data dictionary content always correctly represents definitions of class capabilities and data dictionary translations as used in Heterogeneous Services connections.

**See Also:** *Oracle9i Database Reference* for a description of this parameter.

## Disabling Agent Self-Registration

To disable agent self-registration, set the `HS_AUTOREGISTER` initialization parameter as follows:

```
HS_AUTOREGISTER = FALSE
```

Disabling agent self-registration entails that agent information is not stored in the data dictionary. Consequently, the Heterogeneous Services data dictionary views are not useful sources of information. Nevertheless, the Oracle server still requires information about the class and instance of each agent. If agent self-registration is disabled, the server stores this information in local memory.

## Oracle Database Server SQL Construct Processing

The gateway rewrites SQL statements when the statements need to be translated or post-processed.

For example, consider a program that requests the following from the non-Oracle database system database:

```
SELECT "COL_A" FROM "test"@SYBS
WHERE "COL_A" = INITCAP('jones');
```

The non-Oracle database system database does not recognize `INITCAP`, so the Oracle database server does a table scan of `test` and filters the results locally. The gateway rewrites the `SELECT` statement as follows:

```
SELECT "COL_A" FROM "test"@SYBS
```

The results of the query are sent to the gateway and are filtered by the Oracle database server.

Consider the following `UPDATE` request:

```
UPDATE "test"@SYBS WHERE "COL_A" = INITCAP('jones');
```

In this case, the Oracle database server and the gateway cannot compensate for the lack of support at the non-Oracle database system side, so an error is issued.

If you are performing operations on large amounts of data stored in the non-Oracle database system database, keep in mind that some functions require data to be moved to the integrating Oracle database server before processing can occur.

## Using Synonyms

You can provide complete data location transparency and network transparency by using the synonym feature of the Oracle database server. When a synonym is defined, you do not have to know the underlying table or network protocol. A synonym can be public, which means that all Oracle users can refer to the synonym. A synonym can also be defined as private, which means every Oracle user must have a synonym defined to access the non-Oracle database system table.

The following statement creates a system wide synonym for the EMP table in the schema of user ORACLE in the non-Oracle database system database:

```
SQL> CREATE PUBLIC SYNONYM EMP FOR "ORACLE"."EMP"@SYBS
```

**See Also:** *Oracle9i Database Administrator's Guide* for information about synonyms.

### Example of a Distributed Query

The following example joins data between the Oracle database server, an IBM DB2 database, and the non-Oracle database system database:

```
SQL> SELECT O.CUSTNAME, P.PROJNO, E.ENAME, SUM(E.RATE*P."HOURS")
       FROM ORDERS@DB2 O, EMP@ORACLE9 E, "PROJECTS"@SYBS P
       WHERE O.PROJNO = P."PROJNO"
       AND P."EMPNO" = E.EMPNO
       GROUP BY O.CUSTNAME, P."PROJNO", E.ENAME
```

Through a combination of views and synonyms, using the following SQL statements, the process of distributed queries is transparent to the user:

```
SQL> CREATE SYNONYM ORDERS FOR ORDERS@DB2
SQL> CREATE SYNONYM PROJECTS FOR "PROJECTS"@SYBS
SQL> CREATE VIEW DETAILS (CUSTNAME, PROJNO, ENAME, SPEND)
       AS
       SELECT O.CUSTNAME, P."PROJNO", E.ENAME, SUM(E.RATE*P."HOURS")
       SPEND
       FROM ORDERS O, EMP E, PROJECTS P
       WHERE O.PROJNO = P."PROJNO"
       AND P."EMPNO" = E.EMPNO
       GROUP BY O.CUSTNAME, P."PROJNO", E.ENAME
```

Use the following SQL statement to retrieve information from the data stores in one command:

```
SQL> SELECT * FROM DETAILS;
```

The command retrieves the following table:

CUSTNAME	PROJNO	ENAME	SPEND
-----	-----	-----	-----
ABC Co.	1	Jones	400
ABC Co.	1	Smith	180
XYZ Inc.	2	Jones	400
XYZ Inc.	2	Smith	180



## Copying Data from the Oracle Database Server to the Non-Oracle Database System

Use the SQL\*Plus COPY command to copy data from the local database to the non-Oracle database system database. The syntax is as follows:

```
COPY FROM username/password@db_name
INSERT destination_table USING query
```

The following example selects all rows from the local Oracle EMP table, inserts them into the EMP table on the non-Oracle database system database, and commits the transaction:

```
SQL> COPY FROM SCOTT/TIGER@ORACLE9-
> INSERT SCOTT.EMP@SYBS -
> USING SELECT * FROM EMP
```

The COPY command supports APPEND, CREATE, INSERT, and REPLACE options. However, INSERT is the only option supported when copying to non-Oracle database system. The SQL\*Plus COPY command does not support copying to tables with lowercase table names. Use the following PL/SQL syntax with lowercase table names:

```
DECLARE
    v1 oracle_table.column1%TYPE;
    v2 oracle_table.column2%TYPE;
    v3 oracle_table.column3%TYPE;
    .
    .
    .
    CURSOR cursor_name IS SELECT * FROM oracle_table;
BEGIN
    OPEN cursor_name;
    LOOP
        FETCH cursor_name INTO v1, v2, v3, ... ;
        EXIT WHEN cursor_name%NOTFOUND;
        INSERT INTO destination_table VALUES (v1, v2, v3, ...);
    END LOOP;

    CLOSE cursor_name;
END;
```

**See Also:** *SQL\*Plus User's Guide and Reference* for more information about the COPY command.

The following Oracle SQL INSERT statement is not supported for copying data from the Oracle database server to non-Oracle database system:

```
INSERT INTO table_name SELECT column_list FROM table_name
```

For example, consider the following statement:

```
SQL> INSERT INTO SYBS_TABLE SELECT * FROM MY_LOCAL_TABLE
```

The statement returns the following error message:

```
ORA-2025: All tables in the SQL statement must be at the remote database
```

## Copying Data from the Non-Oracle Database System to the Oracle Database Server

The `CREATE TABLE` command lets you copy data from a non-Oracle database system database to the Oracle database server. To create a table on the local database and insert rows from the non-Oracle database system table, use the following syntax:

```
CREATE TABLE table_name AS query
```

The following example creates the table `EMP` in the local Oracle database and inserts the rows from the `EMP` table of the non-Oracle database system database:

```
SQL> CREATE TABLE EMP AS SELECT * FROM SCOTT."EMP"@SYBS
```

Alternatively, you can use the SQL\*Plus `COPY` command to copy data from the non-Oracle database system database to the Oracle database server.

**See Also:** *SQL\*Plus User's Guide and Reference* for more information about the `COPY` command.

## Heterogeneous Services Data Dictionary Views

You can use the Heterogeneous Services data dictionary views to access information about Heterogeneous Services. This section addresses the following topics:

- [Understanding the Types of Views](#)
- [Understanding the Sources of Data Dictionary Information](#)
- [Using the General Views](#)
- [Using the Transaction Service Views](#)
- [Using the SQL Service Views](#)

### Understanding the Types of Views

The Heterogeneous Services data dictionary views, which all begin with the prefix `HS_`, can be divided into four main types:

- General views
- Views used for the transaction service
- Views used for the SQL service

Most of the data dictionary views are defined for both classes and instances. Consequently, for most types of data there is a `*_CLASS` and an `*_INST` view.

**Table 4–1 Data Dictionary Views for Heterogeneous Services**

View	Type	Identifies
HS_BASE_CAPS	SQL service	All capabilities supported by Heterogeneous Services
HS_BASE_DD	SQL service	All data dictionary translation table names supported by Heterogeneous Services
HS_CLASS_CAPS	Transaction service, SQL service	Capabilities for each class
HS_CLASS_DD	SQL service	Data dictionary translations for each class
HS_CLASS_INIT	General	Initialization parameters for each class
HS_FDS_CLASS	General	Classes accessible from this Oracle9i server
HS_FDS_INST	General	Instances accessible from this Oracle9i server

Like all Oracle data dictionary tables, these views are read-only. Do not use SQL to change the content of any of the underlying tables. To make changes to any of the underlying tables, use the procedures available in the `DBMS_HS` package.

## Understanding the Sources of Data Dictionary Information

The values used for data dictionary content in any particular connection on a Heterogeneous Services database link can come from any of the following sources, in order of precedence:

- Instance information uploaded by the connected Heterogeneous Services agent at the start of the session. This information overrides corresponding content in the Oracle data dictionary, but is never stored into the Oracle data dictionary.
- Instance information stored in the Oracle data dictionary. This data overrides any corresponding content for the connected class.
- Class information stored in the Oracle data dictionary.

If the Oracle database server runs with the `HS_AUTOREGISTER` server initialization parameter set to `FALSE`, then no information is stored automatically in the Oracle data dictionary. The equivalent data is uploaded by the Heterogeneous Services

agent on a connection-specific basis each time a connection is made, with any instance-specific information taking precedence over class information.

---

---

**Note:** It is not possible to determine positively what capabilities and what data dictionary translations are in use for a given session due to the possibility that an agent can upload instance information.

---

---

You can determine the values of Heterogeneous Services initialization parameters by querying the `VALUE` column of the `V$HS_PARAMETER` view. Note that the `VALUE` column of `V$HS_PARAMETER` truncates the actual initialization parameter value from a maximum of 255 characters to a maximum of 64 characters, and it truncates the parameter name from a maximum of 64 characters to a maximum of 30 characters.

### Using the General Views

The views that are common for all services are as follows:

View	Contains
HS_FDS_CLASS HS_FDS_INST	Names of the instances and classes that are uploaded into the Oracle8i data dictionary
HS_CLASS_INIT	Information about the Heterogeneous Services initialization parameters

For example, you can access multiple Sybase gateways from an Oracle database server. After accessing the gateways for the first time, the information uploaded into the Oracle database server could appear as follows:

```
SQL> SELECT * FROM hs_fds_class;
```

FDS_CLASS_NAME	FDS_CLASS_COMMENTS	FDS_CLASS_ID
-----	-----	-----
Sybase816	Uses Sybase driver, R1.1	1
Sybase817	Uses Sybase driver, R1.2	21

Two classes are uploaded: a class that accesses Sybase816 and a class that accesses Sybase817. The data dictionary in the Oracle database server now contains

capability information, SQL translations, and data dictionary translations for both Sybase816 and Sybase817.

In addition to this information, the Oracle database server data dictionary also contains instance information in the `HS_FDS_INST` view for each non-Oracle system instance that is accessed.

## Using the Transaction Service Views

When a non-Oracle system is involved in a distributed transaction, the transaction capabilities of the non-Oracle system and the agent control whether it can participate in distributed transactions. Transaction capabilities are stored in the `HS_CLASS_CAPS` tables.

The ability of the non-Oracle system and agent to support two-phase commit protocols is specified by the 2PC type capability, which can specify one of the following five types.

Read-only (RO)	The non-Oracle system can only be queried with SQL <code>SELECT</code> statements. Procedure calls are not allowed because procedure calls are assumed to write data.
Single-Site (SS)	The non-Oracle system can handle remote transactions but not distributed transactions. That is, it can not participate in the two-phase commit protocol.
Commit Confirm (CC)	The non-Oracle system can participate in distributed transactions. It can participate in the server's two-phase commit protocol but only as the Commit Point Site. That is, it can <i>not</i> prepare data, but it can remember the outcome of a particular transaction if asked by the global coordinator.
Two-Phase Commit	The non-Oracle system can participate in distributed transactions. It can participate in the server's two-phase commit protocol, as a regular two-phase commit node, but not as a Commit Point Site. That is, it can prepare data, but it can <i>not</i> remember the outcome of a particular transaction if asked to by the global coordinator.
Two-Phase Commit Confirm	The non-Oracle system can participate in distributed transactions. It can participate in the server's two-phase commit protocol as a regular two-phase commit node or as the Commit Point Site. That is, it can prepare data and it can remember the outcome of a particular transaction if asked by the global coordinator.

The transaction model supported by the driver and non-Oracle system can be queried from Heterogeneous Services' data dictionary view `HS_CLASS_CAPS`.

One of the capabilities is of the 2PC type:

```
SELECT cap_description, translation
FROM   hs_class_caps
WHERE  cap_description LIKE '2PC%'
AND    fds_class_name='MegaBase6';
```

CAP_DESCRIPTION	TRANSLATION
-----	-----
2PC type (RO-SS-CC-PREP/2P-2PCC)	CC

When the non-Oracle system and agent support distributed transactions, the non-Oracle system is treated like any other Oracle9i server. When a failure occurs during the two-phase commit protocol, the transaction is recovered automatically. If the failure persists, the in-doubt transaction may need to be manually overridden by the database administrator.

### Using the SQL Service Views

Data dictionary views that are specific for the SQL service contain information about:

- SQL capabilities and SQL translations of the non-Oracle data source
- Data Dictionary translations to map Oracle data dictionary views to the data dictionary of the non-Oracle system.

---

---

**Note:** This section describes only a portion of the SQL Service-related capabilities. Because you should never need to alter these settings for administrative purposes, these capabilities are not discussed here.

---

---

### Using Views for Capabilities and Translations

The HS\_\*\_CAPS data dictionary tables contain information about the SQL capabilities of the non-Oracle data source and required SQL translations. These views specify whether the non-Oracle data store or the Oracle database server implements certain SQL language features. If a capability is turned off, then Oracle9i does not send any SQL statements to the non-Oracle data source that require this particular capability, but it still performs post-processing.

### Using Views for Data Dictionary Translations

In order to make the non-Oracle system appear similar to an Oracle database server, Heterogeneous Services connections map a limited set of Oracle data dictionary



views onto the non-Oracle system's data dictionary. This mapping permits applications to issue queries as if these views belonged to an Oracle data dictionary. Data dictionary translations make this access possible. These translations are stored in Heterogeneous Services views whose names are suffixed with `_DD`.

For example, the following `SELECT` statement transforms into a Sybase query that retrieves information about `EMP` tables from the Sybase data dictionary table:

```
SELECT * FROM USER_TABLES@salesdb
WHERE UPPER(TABLE_NAME)='EMP' ;
```

Data dictionary tables can be mimicked instead of translated. If a data dictionary translation is not possible because the non-Oracle data source does not have the required information in its data dictionary, Heterogeneous Services causes it to appear as if the data dictionary table is available, but the table contains no information.

To retrieve information for which Oracle data dictionary views or tables are translated or mimicked for the non-Oracle system, you can issue the following query on the `HS_CLASS_DD` view:

```
SELECT DD_TABLE_NAME, TRANSLATION_TYPE
FROM   HS_CLASS_DD
WHERE  FDS_CLASS_NAME='Sybase' ;
```

DD_TABLE_NAME	TRANSLATION_TYPE
-----	-
ALL_ARGUMENTS	M
ALL_CATALOG	T
ALL_CLUSTERS	T
ALL_CLUSTER_HASH_EXPRESSIONS	M
ALL_COLL_TYPES	M
ALL_COL_COMMENTS	T
ALL_COL_PRIVS	M
ALL_COL_PRIVS_MADE	M
ALL_COL_PRIVS_RECD	M
...	

The translation type `T` specifies that a translation exists. When the translation type is `M`, the data dictionary table is mimicked.

**See Also:** [Appendix D, "Data Dictionary Translation Support"](#) for a list of data dictionary views that are supported through Heterogeneous Services mapping.

## Using the Heterogeneous Services Dynamic Performance Views

The Oracle database server stores information about agents, sessions, and parameter. You can use the `V$` dynamic performance views to access this information. This section contains the following topics:

- [Determining Which Agents Are Running on a Host](#)
- [Determining the Open Heterogeneous Services Sessions](#)

### Determining Which Agents Are Running on a Host

The following view shows generation information about agents:

View	Purpose
<code>V\$HS_AGENT</code>	Identifies the set of Heterogeneous Services agents currently running on a given host, using one row per agent process.

Use this view to determine general information about the agents running on a specified host. The following table shows the most relevant columns (for a description of all the columns in the view, see *Oracle9i Database Reference*):

**Table 4-2** `V$HS_AGENT`

Column	Description
<code>AGENT_ID</code>	Oracle Net session identifier used for connections to agent ( <code>listener.ora SID</code> )
<code>MACHINE</code>	Operating system machine name
<code>PROGRAM</code>	Program name of agent
<code>AGENT_TYPE</code>	Type of agent
<code>FDS_CLASS_ID</code>	The ID of the foreign data store class
<code>FDS_INST_ID</code>	The instance name of the foreign data store

### Determining the Open Heterogeneous Services Sessions

The following view shows which Heterogeneous Services sessions are open for the Oracle database server:

View	Purpose
V\$HS_SESSION	Lists the sessions for each agent, specifying the database link used.

The following table shows the most relevant columns (for an account of all the columns in the view, see *Oracle9i Database Reference*):

**Table 4–3 V\$HS\_SESSION**

Column	Description
HS_SESSION_ID	Unique Heterogeneous Services session identifier
AGENT_ID	Oracle Net session identifier used for connections to agent (listener.ora SID)
DB_LINK	Server database link name used to access the agent NULL means that no database link is used (eg, when using external procedures)
DB_LINK_OWNER	Owner of the database link in DB_LINK

## Determining the Heterogeneous Services Parameters

The following view shows which Heterogeneous Services parameters are set in the Oracle database server:

View	Purpose
V\$HS_PARAMETER	Lists Heterogeneous Services parameters and values registered in the Oracle database server.

The following table shows the most relevant columns (for an account of all the columns in the view, see *Oracle9i Database Reference*):

**Table 4–4 V\$HS\_SESSION**

Column	Description
HS_SESSION_ID	Unique Heterogeneous Services session identifier
PARAMETER	The name of the Heterogeneous Services parameter
VALUE	The value of the Heterogeneous Services parameter

Information about the database link that was used for establishing the distributed connection, the startup time, and the set of initialization parameters used for the session is also available.

All of the runtime information is derived from dynamically updated v\$ tables. The Distributed Access Manager has a refresh capability available through the menu and toolbar that allows users to rerun queries if necessary and update the data. When the data is refreshed, the tool verifies that the set of registered agents remains the same. If it is not, the global view is updated.

**See Also:** *Oracle Enterprise Manager Administrator's Guide* and online help for more information about the Distributed Access Manager.

---

---

# Using Multithreaded Agents

This chapter explains what multithreaded agents are, how they contribute to the overall efficiency of a distributed database system, and how to administer multithreaded agents.

This chapter contains the following sections:

- [Concepts](#)
- [Multithreaded Agent Architecture](#)
- [Multithreaded Agent Administration](#)

---

---

**Note:** Even though Heterogeneous Services supports multithreaded agents, this functionality is not necessarily available in all Heterogeneous Services based gateways. Not only must multithreaded agents have generic support, which Heterogeneous Services provides, but support for multithreaded agents must also be added to the driver.

---

---

## Concepts

This section explains how multithreaded agents contribute to the overall efficiency of Heterogeneous Services and Oracle Transparent Gateways.

This section contains the following topics:

- [The Challenge of Dedicated Agent Architecture](#)
- [The Advantage of Multithreading](#)

### The Challenge of Dedicated Agent Architecture

In the architecture of past releases of Heterogeneous Service, agents are started up on a per user-session and per database link basis. When a user session attempts to access a non-Oracle system by means of a particular database link, an agent process is started up that is exclusively dedicated to that user session and that database link. The agent process terminates only when the user-session ends or when the database link is closed. Separate agent processes are started under the following conditions:

- The same user-session uses two different database links to connect to the same non-Oracle system
- Two different user sessions use the same database link to access the same non-Oracle system.

This architecture is simple and straightforward. However, it has the disadvantage of potentially consuming an unnecessarily large amount of system resources.

For example, suppose that there are several thousand user sessions simultaneously accessing the same non-Oracle system. Because an agent process is started up for each one of them, there are also several thousand agent processes running concurrently as well as several thousand connections open to these agent processes. The agent processes are all running regardless of whether each individual agent process is actually active at the moment or not. Because of this, agent processes and open connections can consume an disproportionate amount of system resources without any discernible benefit.

### The Advantage of Multithreading

Usually, only a small percentage of these agent processes are actually active at a given moment. This makes it possible to more efficiently use system resources by using the multithreaded agent feature of Oracle Transparent Gateways. The multithreaded agent uses a pool of shared agent processes. (The number of these shared agent processes is usually considerably less than the number of user sessions.) The tasks requested by the user sessions are put on a queue and are picked up by the first available multithreaded agent.

**Note Also:** For more information about multithreading, see the following:

- *Oracle9i Database Administrator's Guide*
- *Oracle9i SQL Reference*
- *Oracle Net Services Administrator's Guide*
- *Oracle Net Services Reference Guide*

## Multithreaded Agent Architecture

This section describes the architecture of multithreaded agents.

This section contains the following topics:

- [Overview](#)
- [The Monitor Thread](#)
- [Dispatcher Threads](#)
- [Task Threads](#)

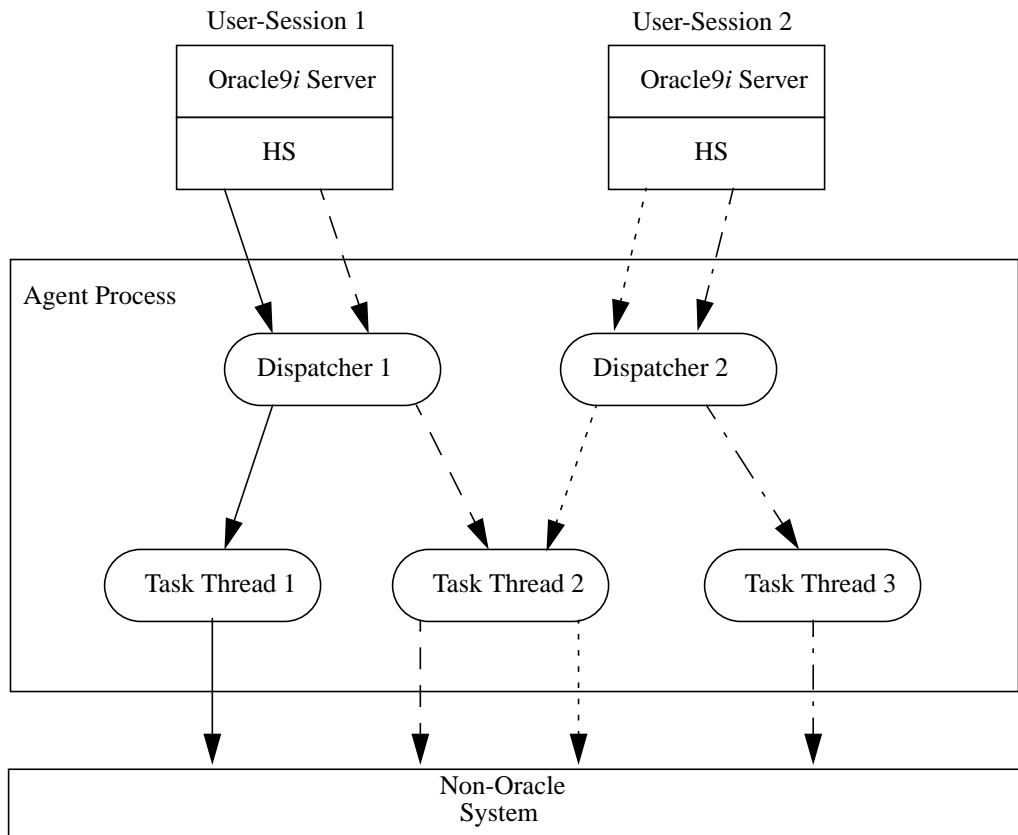
### Overview

In the architecture for multithreaded agents, there are three kinds of threads:

- A single **monitor** thread
- Several **dispatcher** threads
- Several **task** threads.

Typically there are many more task threads than dispatcher threads. The architecture is shown in [Figure 5–1](#).



**Figure 5–1 Multithreaded Agent Architecture**

Each request issued by a user session is represented in [Figure 5–1](#) by a separate type of arrow.

Each request is processed by means of the three different kinds of threads

- The monitor thread is responsible for the following:
  - Maintaining communication with the listener
  - Monitoring the load on the process
  - Starting and stopping threads when required

- The dispatcher threads are responsible for the following:
  - Handling communication with the Oracle server
  - Passing task requests onto the task threads
- The task threads handle requests from the Oracle processes

These three thread types roughly correspond to the Oracle multithreaded server's PMON, dispatcher and shared server processes respectively.

---

---

**Note:** All requests from a user session go through the same dispatcher thread, but can be serviced by different task threads. It is also possible for several task threads to use the same connection to the non-Oracle system.

---

---

Multi-threaded agents are started on a per system-identifier(SID) basis. Each TNS listener that is running on a system listens for incoming connection requests for a set of SIDs. To connect to a process by means of a listener, the SID in the SQL\*Net connect string should be one of the SIDs that the listener is listening for. For each SID, a separate agent process is started and incoming connections for that SID are handed over by the listener to that process.

The agent process is pre-started. A separate agent control utility stops and starts the multithreaded agent itself.

**See Also:** [Multithreaded Agent Administration](#) on page 5-8 for more information on how to start and stop multithreaded agents using the agent control utility.

## The Monitor Thread

The monitor thread is the first thread to be started with an multithreaded agent process. The monitor thread does the following:

- Creates the dispatcher and task threads
- Registers the dispatcher threads it has created with all the listeners that are handling connections to this agent
  - While the dispatcher for this SID is running, the listener does not start a new process when it gets an incoming connection. Instead, the listener hands over the connection to this same dispatcher.
- Monitors the other threads and send load information about the dispatcher threads to all the listener processes handling connections to this agent

- This way, the listeners can hand over incoming connections to the least loaded dispatcher.
- Monitors each of the threads it has created

## Dispatcher Threads

Dispatcher threads do the following:

- Accept incoming connections and task requests from Oracle servers
- Place incoming requests on a queue for a task thread to pick up
- Send results of a request back to the server that issued the request

---

---

**Note:** Once a user session establishes a connection with a dispatcher, all requests from that user-session will go to the same dispatcher until the end of the user session.

---

---

## Task Threads

Task threads do the following:

- Pick up requests from a queue
- Perform the necessary operations.
- Place the results on a queue for a dispatcher to pick up

## Multithreaded Agent Administration

This section explains how you can administer multithreaded agents.

This section contains the following topics.

- [Overview](#)
- [Single Command Mode Commands](#)
- [Shell Mode Commands](#)

### Overview

The multithreaded agent is started and stopped by an agent control utility called `agtctl`, which works much like `lsnrctl`. The main differences are that `lsnrctl` reads a configuration file whereas `agtctl` takes information from the command line and writes it to a control file. There is no equivalent for `listener.ora` as far as `agtctl` is concerned.

You can run `agtctl` in one of two ways:

1. Commands can be run from the UNIX (or DOS) shell  
This mode is called single command mode
2. You can type `'agtctl'` and you will get an `AGTCTL>` prompt and you can type commands from within the `agtctl` shell.

This mode is called shell mode

### Single Command Mode Commands

The commands (in single command mode are) are as follows:

1. Startup

```
agtctl startup agent_name agent_sid
```

2. Shutdown

There are three variants of the shutdown command

1. `agtctl shutdown <sid>`
2. `agtctl shutdown immediate <sid>`
3. `agtctl shutdown abort <sid>`

If you issue the first variant, `agtctl` will talk to the agent and ask it to terminate itself gracefully. In other words, all sessions will complete the operations they are currently doing and then shutdown.

If you issue the second variant, `agtctl` will talk to the agent and tell it terminate immediately. In other words, the agent process will exit immediately regardless of the state of current sessions.

If you issue the third variant, `agtctl` will not talk to the agent at all. It will just issue a system call to kill the agent process.

### 3. Setting parameters

```
agtctl set parameter_value agent_sid
```

### 4. Unsetting parameters

```
agtctl unset parameter agent_sid
```

### 5. Examining parameter values

```
agtctl show parameter agent_sid
```

### 6. Deleting all settings for a particular agent system identifier

```
agtctl delete agent_sid
```

## Shell Mode Commands

In shell mode, you can start `agtctl` by typing `'agtctl'` whereupon you will get an `'AGTCTL>'` prompt.

First, set the name of the agent sid that you are working with by typing

```
set agent_sid agent sid
```

All commands issued after this are assumed to be for this particular sid until the `agent_sid` value is changed.

The commands are all the same as those for the single command mode commands with the exception that you can drop the `'agtctl'` and `agent_sid`.

To set an initialization parameter value, type:

```
set parameter value
```

Use the following table to set your initialization parameters.

**Table 5–1 Initialization Parameters for *agtctl***

<b>parameter</b>	<b>description</b>
<code>max_dispatchers</code>	(maximum number of dispatchers)
<code>tcp_dispatcher</code>	(number of dispatchers listening on tcp - the rest are using ipc).
<code>max_task_threads</code>	(number of task threads)
<code>listener_address</code>	(address on which the listener is listening - needed for registration)
<code>shutdown_address</code>	(address on which the agent should listen for shutdown messages from <i>agtctl</i> )
<code>language</code>	(language name)

---

## Performance Tips

This chapter explains how to optimize distributed SQL statements, how to use partition views with Oracle Transparent Gateways, and how to optimize the performance of distributed queries.

This chapter includes the following sections:

- [Optimizing Heterogeneous Distributed SQL Statements](#)
- [Using Gateways and Partition Views](#)
- [Optimizing Performance of Distributed Queries](#)

## Optimizing Heterogeneous Distributed SQL Statements

When a SQL statement accesses data from non-Oracle systems, it is said to be a heterogeneous distributed SQL statement. To optimize heterogeneous distributed SQL statements, follow the same guidelines as for optimizing distributed SQL statements that access Oracle databases only. However, you must consider that the non-Oracle system usually does not support all the functions and operators that Oracle9i supports.

The Transparent Gateways tell Oracle (at connect time) which functions and operators they do support. If the other data source does not support a function or operator, then Oracle performs that function or operator. In this case, Oracle obtains the data from the other data source and applies the function or operator locally. This affects the way in which the SQL statements are decomposed and can affect performance, especially if Oracle is not on the same machine as the other data source.

## Using Gateways and Partition Views

You can use partition views with Oracle Transparent Gateways release 8 or higher. Make sure you adhere to the following rules:

The cost-based optimizer must be used, by using hints or setting the parameter `OPTIMIZER_MODE` to `ALL_ROWS` or `FIRST_ROWS_K`, or `FIRST_ROWS`.

Indexes used for each partition must be the same. See the gateway-specific documentation to find out whether the gateway sends index information of the non-Oracle system to the Oracle Server. If the gateway sends index information to the optimizer, then make sure that each partition uses the same number of indexes and that you have indexed the same columns. If the gateway does not send index information, then the Oracle optimizer is not aware of the indexes on partitions. Indexes are, therefore, considered to be the same for each partition in the non-Oracle system. If one partition resides on an Oracle server, then you cannot have an index defined on that partition.

The column names and column data types for all branches in the `UNION ALL` view must be the same. Non-Oracle system data types are mapped onto Oracle data types. Make sure that the data types of each partition that reside in the different non-Oracle systems all map to the same Oracle data types. To see how data types are mapped onto Oracle data types, execute a `DESCRIBE` statement in SQL\*Plus.



## Optimizing Performance of Distributed Queries

You can improve performance of distributed queries in several ways:

### **Choose the best SQL statement.**

In many cases, there are several SQL statements that can achieve the same result. If all tables are on the same database, then the difference in performance between these SQL statements might be minimal. But, if the tables are located on different databases, then the difference in performance might be more significant.

### **Use the cost-based optimizer.**

The cost-based optimizer uses indexes on remote tables, considers more execution plans than the rule-based optimizer, and generally gives better results. With the cost-based optimizer, performance of distributed queries is generally satisfactory. Only on rare occasions is it necessary to change SQL statements, create views, or use procedural code.

### **Use views.**

In some situations, views can be used to improve performance of distributed queries. For example:

- Joining several remote tables on the remote database.
- Sending a different table through the network.
- Using procedural code.

On some rare occasions, it can be more efficient to replace a distributed query by procedural code, such as a PL/SQL procedure or a precompiler program. This option is mentioned here only for completeness, not because it is often needed.



---

---

# Generic Connectivity

This chapter describes the configuration and usage of generic connectivity agents.

This chapter contains these topics:

- [What Is Generic Connectivity?](#)
- [Supported Oracle SQL Statements](#)
- [Configuring Generic Connectivity Agents](#)
- [ODBC Connectivity Requirements](#)
- [OLE DB \(SQL\) Connectivity Requirements](#)
- [OLE DB \(FS\) Connectivity Requirements](#)

## What Is Generic Connectivity?

Generic connectivity is intended for low-end data integration solutions requiring the ad hoc query capability to connect from an Oracle database server to non-Oracle database systems. Generic connectivity is enabled by Oracle Heterogeneous Services, allowing you to connect to non-Oracle systems with improved performance and throughput.

Generic connectivity is implemented as either a Heterogeneous Services ODBC agent or a Heterogeneous Services OLE DB agent. An ODBC agent and OLE DB agent are included as part of your Oracle system. Be sure to use the agents shipped with your particular Oracle system, installed in the same `$ORACLE_HOME`.

Any data source compatible with the ODBC or OLE DB standards described in this chapter can be accessed using a generic connectivity agent.

This section contains the following topics:

- [Types of Agents](#)
- [Generic Connectivity Architecture](#)
- [SQL Execution](#)
- [Data Type Mapping](#)
- [Generic Connectivity Restrictions](#)

## Types of Agents

Generic connectivity is implemented as one of the following types of Heterogeneous Services agents:

- ODBC agent for accessing ODBC data providers
- OLE DB agent for accessing OLE DB data providers that support SQL processing—sometimes referred to as **OLE DB (SQL)**
- OLE DB agent for accessing OLE DB data providers without SQL processing support—sometimes referred to as **OLE DB (FS)**

Each user session receives its own dedicated agent process spawned by the first use in that user session of the database link to the non-Oracle system. The agent process ends when the user session ends.

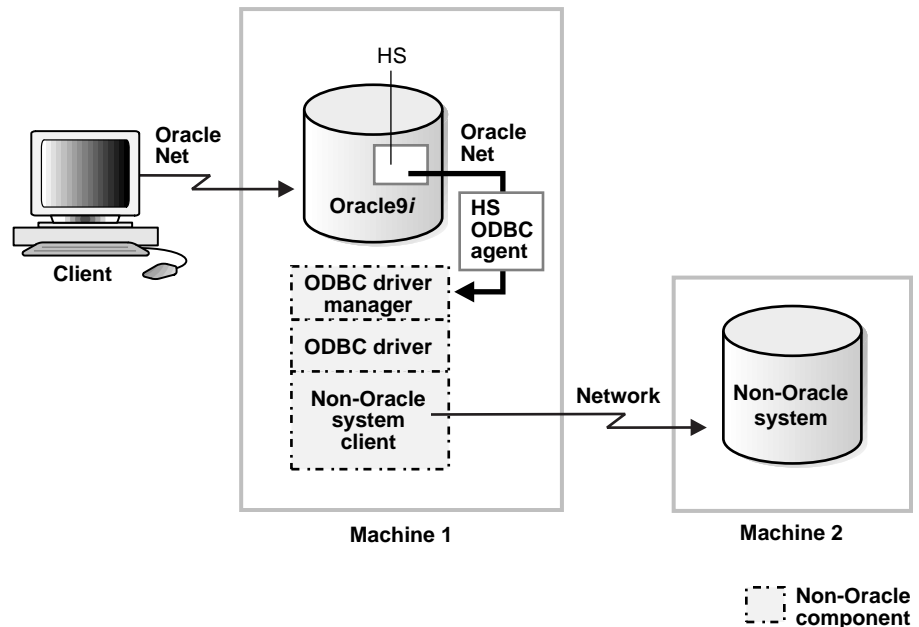
## Generic Connectivity Architecture

To access the non-Oracle data store using generic connectivity, the agents work with an ODBC or OLE DB driver. The Oracle database server provides support for the ODBC or OLE DB driver interface. The driver that you use must be on the same platform as the agent. The non-Oracle data stores can reside on the same machine as the Oracle database server or on a different machine.

### Oracle and Non-Oracle Systems on Separate Machines

Figure 7-1 shows an example of a configuration in which an Oracle and non-Oracle database are on separate machines, communicating through an Heterogeneous Services ODBC agent.

*Figure 7-1 Oracle and Non-Oracle Systems on a Separate Machines*



In this configuration:


1. A client connects to the Oracle database server through Oracle Net

2. The Heterogeneous Services component of the Oracle database server connects through Oracle Net to the Heterogeneous Services ODBC agent
3. The agent communicates with the following non-Oracle components:
  - An ODBC driver manager
  - An ODBC driver
  - A non-Oracle client application

This client connects to the non-Oracle data store through a network.

### **Oracle and Non-Oracle Systems on the Same Machine**

[Figure 7-2](#) shows an example of a different configuration in which an Oracle and non-Oracle database are on the same machine, again communicating through an Heterogeneous Services ODBC agent.

**Figure 7–2 Oracle and non-Oracle Systems on the Same Machine**

In this configuration:

1. A client connects to the Oracle database server through Oracle Net
2. The Heterogeneous Services component of the Oracle database server connects through Oracle Net to the Heterogeneous Services ODBC agent
3. The agent communicates with the following non-Oracle components:
  - An ODBC driver manager
  - An ODBC driver

The driver then allows access to the non-Oracle data store.

---

---

**Note:** The ODBC driver may require non-Oracle client libraries even if the non-Oracle database is located on the same machine.

---

---

## SQL Execution

SQL statements sent using a generic connectivity agent are executed differently depending on the type of agent you are using: ODBC, OLE DB (SQL), or OLE DB (FS). For example, if a SQL statement involving tables is sent using an ODBC agent for a file-based storage system, the file can be manipulated as if it were a table in a relational database. The naming conventions used at the non-Oracle system can also depend on whether you are using an ODBC or OLE DB agent.

## Data Type Mapping

The Oracle database server maps the data types used in ODBC and OLE DB compliant data sources to supported Oracle data types. When the results of a query are returned, the Oracle database server converts the ODBC or OLE DB data types to Oracle data types. For example, the ODBC data type `SQL_TIMESTAMP` and the OLE DB data type `DBTYPE_DBTIMESTAMP` are converted to Oracle's `DATE` data type.

## Generic Connectivity Restrictions

Generic connectivity restrictions include:

- A table including a BLOB column must have a separate column that serves as a primary key
- BLOB/CLOB data cannot be read through passthrough queries
- Updates or deletes that include unsupported functions within a `WHERE` clause are not allowed
- Stored procedures are not supported
- Generic connectivity agents cannot participate in distributed transactions; they support single-site transactions only



## Supported Oracle SQL Statements

Generic connectivity supports the following statements, but only if the ODBC or OLE DB driver and non-Oracle system can execute them *and* the statements contain supported Oracle SQL functions:

- DELETE
- INSERT
- SELECT
- UPDATE

Only a limited set of functions are assumed to be supported by the non-Oracle system. Most Oracle functions have no equivalent function in this limited set. Consequently, although post-processing is performed by the Oracle database server, many Oracle functions are not supported by generic connectivity, possibly impacting performance.

If an Oracle SQL function is not supported by generic connectivity, then this function is not supported in DELETE, INSERT, or UPDATE statements. In SELECT statements, these functions are evaluated by the Oracle database server and post-processed after they are returned from the non-Oracle system.

If an unsupported function is used in a DELETE, INSERT, or UPDATE statement, it generates this Oracle error:

```
ORA-02070: database db_link_name does not support function in this context
```

## Functions Supported by Generic Connectivity

Generic connectivity assumes that the following minimum set of SQL functions is supported:

- AVG(*exp*)
- LIKE(*exp*)
- COUNT(\*)
- MAX(*exp*)
- MIN(*exp*)
- NOT

## Configuring Generic Connectivity Agents

To implement generic connectivity on a non-Oracle data source, you must set the agent parameters.

This section contains the following topics:

- [Creating the Initialization File](#)
- [Editing the Initialization File](#)
- [Setting Initialization Parameters for an ODBC-based Data Source](#)
- [Setting Initialization Parameters for an OLE DB-based Data Source](#)

### Creating the Initialization File

You must create and customize an initialization file for your generic connectivity agent. Oracle Corporation supplies sample initialization files named `inithsagent.ora`, where *agent* is `odbc` or `oledb`, indicating which agent the sample file can be used for, as in the following:

```
inithsodbc.ora  
inithsoledb.ora
```

The sample files are stored in the `$ORACLE_HOME/hs/admin` directory.

To create an initialization file for an ODBC or OLE DB agent, copy the applicable sample initialization file and rename the file to `initHS_SID.ora`, where *HS\_SID* is the system identifier you want to use for the instance of the non-Oracle system to which the agent connects.

The *HS\_SID* is also used to identify how to connect to the agent when you configure the listener by modifying the `listener.ora` file. The *HS\_SID* you add to the `listener.ora` file must match the *HS\_SID* in an `initHS_SID.ora` file, because the agent spawned by the listener searches for a matching `initHS_SID.ora` file. That is how each agent process gets its initialization information. When you copy and rename your `initHS_SID.ora` file, ensure it remains in the `$ORACLE_HOME/hs/admin` directory.

### Editing the Initialization File

Customize the `initHS_SID.ora` file by setting the parameter values used for generic connectivity agents to values appropriate for your system, agent, and drivers. You must edit the `initHS_SID.ora` file to change the `HS_FDS_CONNECT_`

INFO initialization parameter. `HS_FDS_CONNECT_INFO` specifies the information required for connecting to the non-Oracle system.

**See Also:** "[Initialization Parameters](#)" on page 4-6 for more information on parameters.

Set the parameter values as follows:

```
[SET][PRIVATE] parameter=value
```

where:

`[SET][PRIVATE]` are optional keywords. If you do not specify either `SET` or `PRIVATE`, the parameter and value are simply used as an initialization parameter for the agent.

`SET` specifies that in addition to being used as an initialization parameter, the parameter value is set as an environment variable for the agent process.

`PRIVATE` specifies that the parameter value is private and not transferred to the Oracle database server and does not appear in `V$` tables or in an graphical user interfaces.

`SET PRIVATE` specifies that the parameter value is set as an environment variable for the agent process and is also private (not transferred to the Oracle database server, not appearing in `V$` tables or graphical user interfaces).

*parameter* is the Heterogeneous Services initialization parameter that you are specifying. See "[Initialization Parameters](#)" on page 4-6 for a description of all Heterogeneous Services parameters and their possible values. The parameter is case-sensitive.

*value* is the value you want to specify for the Heterogeneous Services parameter. The value is case-sensitive.

For example, to enable tracing for an agent, set the `HS_FDS_TRACE_LEVEL` parameter as follows:

```
HS_FDS_TRACE_LEVEL=ON
```

Typically, most parameters are only needed as initialization parameters, so you do not need to use `SET` or `PRIVATE`. Use `SET` for parameter values that the drivers or non-Oracle system need as environment variables.

PRIVATE is only supported for the follow Heterogeneous Services parameters:

- HS\_FDS\_CONNECT\_INFO
- HS\_FDS\_SHAREABLE\_NAME
- HS\_FDS\_TRACE\_LEVEL
- HS\_FDS\_TRACE\_FILE\_NAME

You should only use PRIVATE for these parameters if the parameter value includes sensitive information such as a username or password.

### Setting Initialization Parameters for an ODBC-based Data Source

The settings for the initialization parameters vary depending on the type of operating system.

#### Setting Agent Parameters on Windows NT

Specify a File data source name (DSN) or a System DSN which has previously been defined using the ODBC Driver Manager.

When connecting using a File DSN, specify the value as follows:

```
HS_FDS_CONNECT_INFO=FILEDSN=file_dsn
```

When connecting using a System DSN, specify the value as follows:

```
HS_FDS_CONNECT_INFO=system_dsn
```

If you are connecting to the data source through the driver for that data source, precede the DSN by the name of the driver, followed by a semi-colon (;).

**Setting Parameters on NT: Example** Assume a System DSN has been defined in the Windows ODBC Data Source Administrator. In order to connect to this SQL Server database through the gateway, the following line is required in `initHS_SID.ora`:

```
HS_FDS_CONNECT_INFO=sqlserver7
```

where `sqlserver7` is the name of the System DSN defined in the Windows ODBC Data Source Administrator.

The following procedure enables you to define a System DSN in the Windows ODBC Data Source Administrator:

1. From the **Start** menu, choose **Settings > Control Panel** and select the **ODBC** icon.

2. Select the **System DSN** tab to display the system data sources.
3. Click **Add**.
4. From the list of installed ODBC drivers, select the name of the driver that the data source will use. For example, select **SQL Server**.
5. Click **Finish**.
6. Enter a name for the DSN and an optional description. Enter other information depending on the ODBC driver. For example, for SQL Server enter the SQL Server machine.

---

**Note:** The name entered for the DSN must match the value of the parameter `HS_FDS_CONNECT_INFO` that is specified in `initHS_SID.ora`.

---

7. Continue clicking **Next** and answering the prompts until you click **Finish**.
8. Click **OK** until you exit the ODBC Data Source Administrator.

### Setting Agent Parameters on UNIX platforms

Specify a DSN and the path of the ODBC shareable library, as follows:

```
HS_FDS_CONNECT_INFO=dsn_value
HS_FDS_SHAREABLE_NAME=full_odbc_library_path_of_odbc_driver
```

`HS_FDS_CONNECT_INFO` is required for all platforms for an ODBC agent. `HS_FDS_SHAREABLE_NAME` is required on UNIX platforms for an ODBC agent. Other initialization parameters have defaults or are optional. You can use the default values and omit the optional parameters, or you can specify the parameters with values tailored for your installation.

---

**Note:** Before deciding to accept the default values or change them, see "[Initialization Parameters](#)" on page 4-6 for detailed information on all the initialization parameters.

---

**Setting Parameters on UNIX: Example** Assume that the `odbc.ini` file for connecting to Informix using the Intersolve ODBC driver is located in `/opt/odbc` and includes the following information:

```
[ODBC Data Sources]
```

```
Informix=INTERSOLV 3.11 Informix Driver
```

```
[Informix]
Driver=/opt/odbc/lib/ivinf13.so
Description=Informix
Database=personnel@osf_inf72
HostName=osf
LogonID=uid
Password=pwd
```

In order to connect to this Informix database through the gateway, the following lines are required in `initHS_SID.ora`:

```
HS_FDS_CONNECT_INFO=Informix
HS_FDS_SHAREABLE_NAME=/opt/odbc/lib/libodbc.so
set INFORMIXDIR=/users/inf72
set INFORMIXSERVER=osf_inf72
set ODBCINI=/opt/odbc/odbc.ini
```

Note that the set statements are optional as long as they are specified in the working account. Each database has its own set statements.

The `HS_FDS_CONNECT_INFO` parameter value must match the ODBC data source name in the `odbc.ini` file.

## Setting Initialization Parameters for an OLE DB-based Data Source

You can only set these parameters on the Windows NT platform.

Specify a data link (UDL) that has previously been defined:

```
SET/PRIVATE/SET PRIVATE HS_FDS_CONNECT_INFO="UDLFILE=data_link"
```

Or, specify the connection details directly:

```
SET/PRIVATE/SET PRIVATE HS_FDS_CONNECT_INFO="provider;db[,CATALOG=catalog]"
```

where:

*provider* is the name of the provider as it appears in the registry. The value is case sensitive.

*db* is the name of the database

*catalog* is the name of the catalog

---



---

**Note:** If the parameter value includes an equal sign (=), then it must be surrounded by quotation marks.

---



---

HS\_FDS\_CONNECT\_INFO is required for an OLE DB agent. Other initialization parameters have defaults or are optional. You can use the default values and omit the optional parameters, or you can specify the parameters with values tailored for your installation.

---



---

**Note:** Before deciding to accept the default values or change them, see "[Initialization Parameters](#)" on page 4-6 for detailed information on all the initialization parameters.

---



---

## ODBC Connectivity Requirements

To use an ODBC agent, you must have an ODBC driver installed on the same machine as the Oracle database server. On Windows NT, you must have an ODBC driver manager also located on the same machine. The ODBC driver manager and driver must meet the following requirements:

- On Windows NT machines, a thread-safe, 32-bit ODBC driver Version 2.x or 3.x is required. You can use the native driver manager supplied with your Windows NT system.
- On UNIX machines, ODBC driver Version 2.5 is required. A driver manager is not required.

The ODBC driver and driver manager on Windows NT must conform to ODBC application program interface (API) conformance Level 1 or higher. If the ODBC driver or driver manager does not support multiple active ODBC cursors, then it restricts the complexity of SQL statements that you can execute using generic connectivity.

The ODBC driver you use must support all of the core SQL ODBC data types and should support SQL grammar level SQL\_92. The ODBC driver should also expose the following ODBC APIs:

**Table 7-1 ODBC Functions** (Page 1 of 3)

ODBC Function	Comment
SQLAllocConnect	

**Table 7-1 ODBC Functions** (Page 2 of 3)

<b>ODBC Function</b>	<b>Comment</b>
SQLAllocEnv	
SQLAllocStmt	
SQLBindCol	
SQLBindParameter	
SQLColumns	
SQLConnect	
SQLDescribeCol	
SQLDisconnect	
SQLDriverConnect	
SQLError	
SQLExecDirect	
SQLExecute	
SQLExtendedFetch	Recommended if used by the non-Oracle system.
SQLFetch	
SQLForeignKeys	Recommended if used by the non-Oracle system.
SQLFreeConnect	
SQLFreeEnv	
SQLFreeStmt	
SQLGetConnectOption	
SQLGetData	
SQLGetFunctions	
SQLGetInfo	
SQLGetTypeInfo	
SQLNumParams	Recommended if used by the non-Oracle system.
SQLNumResultCols	
SQLParamData	
SQLPrepare	



**Table 7-1 ODBC Functions** (Page 3 of 3)

ODBC Function	Comment
SQLPrimaryKeys	Recommended if used by the non-Oracle system.
SQLProcedureColumns	Recommended if used by the non-Oracle system.
SQLProcedures	Recommended if used by the non-Oracle system.
SQLPutData	
SQLRowCount	
SQLSetConnectOption	
SQLSetStmtOption	
SQLStatistics	
SQLTables	
SQLTransact	Recommended if used by the non-Oracle system.

## OLE DB (SQL) Connectivity Requirements

These requirements apply to OLE DB data providers that have an SQL processing capability and expose the OLE DB interfaces.

Generic connectivity passes the username and password to the provider when calling `IDBInitialize::Initialize()`.

OLE DB (SQL) connectivity requires that the data provider expose the following OLE DB interfaces:

**Table 7-2 OLE DB (SQL) Interfaces**

Interface	Methods
IAccessor	CreateAccessor, ReleaseAccessor
IColumnsInfo	GetColumnsInfo (Command and Rowset objects)
ICommand	Execute
ICommandPrepare	Prepare
ICommandProperties	SetProperties
ICommandText	SetCommandText
ICommandWithParameters	GetParameterInfo

**Table 7–2 OLE DB (SQL) Interfaces**

<b>Interface</b>	<b>Methods</b>
IDBCreateCommand	CreateCommand
IDBCreateSession	CreateSession
IDBInitialize	Initialize
IDBSchemaRowset	GetRowset (tables, columns, indexes; optionally also procedures, procedure parameters)
IErrorInfo <sup>1</sup>	GetDescription, GetSource
IErrorRecords	GetErrorInfo
ILockBytes (OLE) <sup>2</sup>	Flush, ReadAt, SetSize, Stat, WriteAt
IRowset	GetData, GetNextRows, ReleaseRows, RestartPosition
IStream (OLE)	Read, Seek, SetSize, Stat, Write
ISupportErrorInfo	InterfaceSupportsErrorInfo
ITransactionLocal (optional)	StartTransaction, Commit, Abort

<sup>1</sup> You can also use IErrorLookup with the GetErrorDescription method.

<sup>2</sup> Required only if BLOBs are used in the OLE DB provider.

## OLE DB (FS) Connectivity Requirements

These requirements apply to OLE DB data providers that do not have SQL processing capabilities. If the provider exposes them, then OLE DB (FS) connectivity uses OLE DB Index interfaces.

OLE DB (FS) connectivity requires that the data provider expose the following OLE DB interfaces:

**Table 7–3 OLE DB (FS) Interfaces** (Page 1 of 2)

<b>Interface</b>	<b>Methods</b>
IAccessor	CreateAccessor, ReleaseAccessor
IColumnsInfo	GetColumnsInfo (Command and Rowset objects)

**Table 7-3 OLE DB (FS) Interfaces** (Page 2 of 2)

Interface	Methods
IOpenRowset	OpenRowset
IDBCreateSession	CreateSession
IRowsetChange	DeleteRows, SetData, InsertRow
IRowsetLocate	GetRowsByBookmark
IRowsetUpdate	Update (optional)
IDBInitialize	Initialize, Uninitialize
IDBSchemaRowset	GetRowset (tables, columns, indexes; optionally also procedures, procedure parameters)
ILockBytes (OLE) <sup>1</sup>	Flush, ReadAt, SetSize, Stat, WriteAt
IRowsetIndex <sup>2</sup>	SetRange
IErrorInfo <sup>3</sup>	GetDescription, GetSource
IErrorRecords	GetErrorInfo
IRowset	GetData, GetNextRows, ReleaseRows, RestartPosition
IStream (OLE)	Read, Seek, SetSize, Stat, Write
ITransactionLocal (optional)	StartTransaction, Commit, Abort
ISupportErrorInfo	InterfaceSupportsErrorInfo
ITableDefinition	CreateTable, DropTable
IDBProperties	SetProperties

<sup>1</sup> Required only if BLOBs are used in the OLE DB provider.

<sup>2</sup> Required only if indexes are used in the OLE DB provider.

<sup>3</sup> You can use IErrorLookup with the GetErrorDescription method as well.

Because OLE DB (FS) connectivity is generic, it can connect to a number of different data providers that expose OLE DB interfaces. Every such data provider must meet the certain requirements.

---

---

**Note:** The data provider must expose bookmarks. This enables tables to be updated. Without bookmarks being exposed, the tables are read-only.

---

---

## Data Source Properties

The OLE DB data source must support the following initialization properties:

- DBPROP\_INIT\_DATASOURCE
- DBPROP\_AUTH\_USERID

---

---

**Note:** Required if the userid has been supplied in the security file

---

---

- DBPROP\_AUTH\_PASSWORD

---

---

**Note:** Required if the userid and password have been supplied in the security file

---

---

The OLE DB data source must also support the following rowset properties:

- DBPROP\_IRowsetChange = TRUE
- DBPROP\_UPDATABILITY = CHANGE+DELETE+INSERT
- DBPROP\_OWNUPDATEDELETE = TRUE
- DBPROP\_OWNINSERT = TRUE
- DBPROP\_OTHERUPDATEDELETE = TRUE
- DBPROP\_CANSROLLBACKWARDS = TRUE
- DBPROP\_IRowsetLocate = TRUE
- DBPROP\_OTHERINSERT = FALSE

---

# Heterogeneous Services Initialization Parameters

Heterogeneous Services initialization files, like all Oracle parameter files, are configuration settings stored as a text file in

You can set Heterogeneous Services parameters by editing the Oracle Transparent Gateway initialization file, or by using the `DBMS_HS` package to set them in the data dictionary. String values for Heterogeneous Services parameters must be lowercase.

This section contains the following topics:

- [HS\\_COMMIT\\_POINT\\_STRENGTH](#)
- [HS\\_DB\\_DOMAIN](#)
- [HS\\_DB\\_INTERNAL\\_NAME](#)
- [HS\\_DB\\_NAME](#)
- [HS\\_DESCRIBE\\_CACHE\\_HWM](#)
- [HS\\_FDS\\_CONNECT\\_INFO](#)
- [HS\\_FDS\\_SHAREABLE\\_NAME](#)
- [HS\\_FDS\\_TRACE\\_LEVEL](#)
- [HS\\_LANGUAGE](#)
- [HS\\_LONG\\_PIECE\\_TRANSFER\\_SIZE](#)
- [HS\\_NLS\\_DATE\\_FORMAT](#)
- [HS\\_NLS\\_DATE\\_LANGUAGE](#)
- [HS\\_NLS\\_NCHAR](#)
- [HS\\_NLS\\_TIMESTAMP\\_FORMAT](#)

- 
- HS-NLS\_TIMESTAMP\_TZ\_FORMAT
  - HS\_OPEN\_CURSORS
  - HS\_ROWID\_CACHE\_SIZE
  - HS\_RPC\_FETCH\_REBLOCKING
  - HS\_RPC\_FETCH\_SIZE
  - HS\_TIME\_ZONE
  - IFILE

---

## HS\_COMMIT\_POINT\_STRENGTH

---

<b>Default value:</b>	0
<b>Range of values:</b>	0 to 255

---

Specifies a value that determines the commit point site in a heterogeneous distributed transaction. HS\_COMMIT\_POINT\_STRENGTH is similar to COMMIT\_POINT\_STRENGTH, described in the *Oracle9i Database Reference*.

Set HS\_COMMIT\_POINT\_STRENGTH to a value relative to the importance of the site that is the commit point site in a distributed transaction. The Oracle database server or non-Oracle system with the highest commit point strength becomes the commit point site. To ensure that a non-Oracle system *never* becomes the commit point site, set the value of HS\_COMMIT\_POINT\_STRENGTH to zero.

HS\_COMMIT\_POINT\_STRENGTH is important only if the non-Oracle system can participate in the two-phase protocol as a regular two-phase commit partner and as the commit point site. This is only the case if the transaction model is two-phase commit confirm (2PCC).

## HS\_DB\_DOMAIN

---

<b>Default value:</b>	WORLD
<b>Range of values:</b>	1 to 119 characters

---

Specifies a unique network sub-address for a non-Oracle system. HS\_DB\_DOMAIN is similar to DB\_DOMAIN, described in the *Oracle9i Database Administrator's Guide* and the *Oracle9i Database Reference*. HS\_DB\_DOMAIN is required if you use the Oracle Names server. HS\_DB\_NAME and HS\_DB\_DOMAIN define the global name of the non-Oracle system.

---

---

**Note:** HS\_DB\_NAME and HS\_DB\_DOMAIN must combine to form a unique address.

---

---

## HS\_DB\_INTERNAL\_NAME

---

<b>Default value:</b>	01010101
-----------------------	----------

---

---

<b>Range of values:</b>	1 to 16 hexadecimal characters
-------------------------	--------------------------------

---

Specifies a unique hexadecimal number identifying the instance to which the Heterogeneous Services agent is connected. This parameter's value is used as part of a transaction ID when global name services are activated. Specifying a non-unique number can cause problems when two-phase commit recovery actions are necessary for a transaction.

## **HS\_DB\_NAME**

---

<b>Default value:</b>	HO
-----------------------	----

<b>Range of values:</b>	1 to 8 lowercase characters
-------------------------	-----------------------------

---

Specifies a unique alphanumeric name for the data store given to the non-Oracle system. This name identifies the non-Oracle system within the cooperative server environment. `HS_DB_NAME` and `HS_DB_DOMAIN` define the global name of the non-Oracle system.

## **HS\_DESCRIBE\_CACHE\_HWM**

---

<b>Default value:</b>	100
-----------------------	-----

<b>Range of values:</b>	1 to 4000
-------------------------	-----------

---

Specifies the maximum number of entries in the describe cache used by Heterogeneous Services. This limit is known as the describe cache high water mark. The cache contains descriptions of the mapped tables that Heterogeneous Services reuses so that it does not have to re-access the non-Oracle data store.

If you are accessing many mapped tables, then increase the high water mark to improve performance. Note that increasing the high water mark improves performance at the cost of memory usage.

## **HS\_FDS\_CONNECT\_INFO**

---

<b>Default value:</b>	none
-----------------------	------



---

<b>Range of values:</b>	not applicable
-------------------------	----------------

---

Specifies the information needed to bind to the data provider, that is, the non-Oracle system. For generic connectivity, you can bind to an ODBC-based data source or to an OLE DB-based data source. The information that you provide depends on the platform and whether the data source is ODBC or OLE DB-based.

This parameter is required if you are using generic connectivity.

**ODBC-based Data Source on Windows:** You can use either a File DSN or a System DSN as follows:

- When connecting using a File DSN the parameter format is:

```
HS_FDS_CONNECT_INFO=FILEDSN=file_dsn
```

- When connecting using a System DSN the parameter format is:

```
HS_FDS_CONNECT_INFO=system_dsn
```

If you are connecting to the data source through the driver for that data source, then precede the DSN by the name of the driver, followed by a semi-colon (;).

**ODBC-based Data Source on UNIX:** Use a DSN with the following format:

```
HS_FDS_CONNECT_INFO=dsn
```

**OLE DB-based Data Source (Windows NT Only):** Use a universal data link (UDL) with the following formats:

- `HS_FDS_CONNECT_INFO="UDLFILE=data_link"`
- `HS_FDS_CONNECT_INFO="data_link_provider;db[ , CATALOG=catalog]"`

which allows you to specify the connection details directly, and where:

- *data\_link\_provider* is the case-sensitive name of the provider as it appears in the registry
- *db* is the name of the database
- *catalog* is the name of the catalog

---

---

**Note:** Whenever the parameter value includes an equal sign (=), it must be enclosed in quotation marks.

---

---

## HS\_FDS\_SHAREABLE\_NAME

---

<b>Default value:</b>	none
<b>Range of values:</b>	not applicable

---

Specifies the full path name to the ODBC library. This parameter is required when you are using generic connectivity to access data from an ODBC provider on a UNIX machine.

## HS\_FDS\_TRACE\_LEVEL

---

<b>Default value:</b>	OFF
<b>Range of values:</b>	ON or OFF

---

Specifies whether error tracing is enabled or disabled for generic connectivity. Enable the tracing to see which error messages occur when you encounter problems. The results are written to a generic connectivity log file, in the `/log` directory under the `$ORACLE_HOME` directory.

## HS\_LANGUAGE

---

<b>Default value:</b>	System-specific
<b>Range of values:</b>	Any valid language name (up to 255 characters)

---

Provides Heterogeneous Services with character set, language, and territory information of the non-Oracle data source. The value must use the following format:

*language[\_territory.character\_set]*

---

---

**Note:** The national language support initialization parameters affect error messages, the data for the SQL Service, and parameters in distributed external procedures.

---

---

### Character sets

Ideally, the character sets of the Oracle database server and the non-Oracle data source are the same. If they are not the same, Heterogeneous Services attempts to translate the character set of the non-Oracle data source to the Oracle database character set, and back again. The translation can degrade performance. In some cases, Heterogeneous Services cannot translate a character from one character set to another.

---

---

**Note:** The specified character set must be a superset of the operating system character set on the platform where the agent is installed.

---

---

### Language

The language component of the `HS_LANGUAGE` initialization parameter determines:

- Day and month names of dates
- AD, BC, PM, and AM symbols for date and time
- Default sorting mechanism

Note that Oracle9i does not determine the language for error messages for the generic Heterogeneous Services messages (ORA-25000 through ORA-28000). These are controlled by the session settings in the Oracle database server.

---

---

**Note:** Use the `HS-NLS_DATE_LANGUAGE` initialization parameter to set the day and month names, and the AD, BC, PM, and AM symbols for dates and time independently from the language.

---

---

### Territory

The territory clause specifies the conventions for day and week numbering, default date format, decimal character and group separator, and ISO and local currency symbols. Note that:

- 
- You can override the date format using the initialization parameter `HS_NLS_DATE_FORMAT`.
  - The level of National Language Support between the Oracle database server and the non-Oracle data source depends on how the driver is implemented. See the installation documentation for your platform for more information about the level of National Language Support.

## **HS\_LONG\_PIECE\_TRANSFER\_SIZE**

---

<b>Default value:</b>	64 KB
<b>Range of values:</b>	Any value up to 2 GB

---

Sets the size of the piece of `LONG` data being transferred. A smaller piece size means less memory requirement, but more round trips to fetch all the data. A larger piece size means fewer round trips, but more of a memory requirement to store the intermediate pieces internally. Thus, the initialization parameter can be used to tune a system for the best performance, with the best trade-off between round trips and memory requirements.

## **HS\_NLS\_DATE\_FORMAT**

---

<b>Default value:</b>	Value determined by <code>HS_LANGUAGE</code> parameter
<b>Range of values:</b>	Any valid date format mask (up to 255 characters)

---

Defines the date format for dates used by the target system. This parameter has the same function as the `NLS_DATE_FORMAT` parameter for an Oracle database server. The value of can be any valid date mask listed in the *Oracle9i Database Reference*, but must match the date format of the target system. For example, if the target system stores the date February 14, 2001 as `2001/02/14`, set the parameter to `yyyy/mm/dd`. Note that characters must be lowercase.

## **HS\_NLS\_DATE\_LANGUAGE**

---

<b>Default value:</b>	Value determined by <code>HS_LANGUAGE</code> parameter
<b>Range of values:</b>	Any valid <code>NLS_LANGUAGE</code> value (up to 255 characters)

---

---

Specifies the language used in character date values coming from the non-Oracle system. Date formats can be language independent. For example, if the format is dd/mm/yyyy, all three components of the character date are numbers. In the format dd-mon-yyyy, however, the month component is the name abbreviated to three characters. The abbreviation is very much language dependent. For example, the abbreviation for the month April is "apr", which in French is "avr" (Avril).

Heterogeneous Services assumes that character date values fetched from the non-Oracle system are in this format. Also, Heterogeneous Services sends character date bind values in this format to the non-Oracle system.

## HS\_NLS\_NCHAR

---

<b>Default value:</b>	Value determined by <code>HS_LANGUAGE</code> parameter
<b>Range of values:</b>	Any valid national character set (up to 255 characters)

---

Informs Heterogeneous Services of the value of the national character set of the non-Oracle data source. This value is the non-Oracle equivalent to the `NATIONAL CHARACTER SET` parameter setting in the Oracle `CREATE DATABASE` statement. The `HS_NLS_NCHAR` value should be the character set ID of a character set supported by the Oracle `NLSRTL` library.

**See Also:** [HS\\_LANGUAGE](#) on page A-6.

## HS\_NLS\_TIMESTAMP\_FORMAT

---

<b>Default value:</b>	Derived from <code>NLS_TERRITORY</code>
<b>Range of values:</b>	Any valid datetime format mask

---

Defines the timestamp format for dates used by the target system. This parameter has the same function as the `NLS_TIMESTAMP_FORMAT` parameter for an Oracle database server. The value of can be any valid timestamp mask listed in the Oracle9i Database Reference, but it must match the date format of the target system. Note that characters must be lowercase. For example:

```
HS_NLS_TIMESTAMP_FORMAT = yyyy-mm-dd hh:mi:ss.ff
```

---

## HS\_NLS\_TIMESTAMP\_TZ\_FORMAT

---

**Default value:** Dynamic. Scope= ALTER SESSION  
NLS\_TIMESTAMP\_TZ\_FORMAT

**Range of values:** Derived from NLS\_TERRITORY

---

Defines the default timestamp with time zone format for the timestamp with time zone format used by the target system. This parameter has the same function as the NLS\_TIMESTAMP\_TZ\_FORMAT parameter for an Oracle database server. The value of can be any valid timestamp with time zone mask listed in the Oracle9i Database Reference, but must match the date format of the target system. Note that characters must be lowercase. For example:

```
HS_NLS_TIMESTAMP_TZ_FORMAT = yyyy-mm-dd hh:mi:ss.ff tzh:tzm
```

## HS\_OPEN\_CURSORS

---

**Default value:** 50

**Range of values:** 1 - value of Oracle's OPEN\_CURSORS initialization parameter

---

Defines the maximum number of cursors that can be open on one connection to a non-Oracle system instance.

The value never exceeds the number of open cursors in the Oracle database server. Therefore, setting the same value as the OPEN\_CURSORS initialization parameter in the Oracle database server is recommended.

## HS\_ROWID\_CACHE\_SIZE

---

**Default value:** 3

**Range of values:** 1 to 32767

---

Specifies the size of the Heterogeneous Services cache containing the non-Oracle system equivalent of ROWIDs. The cache contains non-Oracle system ROWIDs needed to support the WHERE CURRENT OF clause in a SQL statement or a SELECT FOR UPDATE statement.

---

When the cache is full, the first slot in the cache is reused, then the second, and so on. Only the last `HS_ROWID_CACHE_SIZE` non-Oracle system ROWIDs are cached.

## HS\_RPC\_FETCH\_REBLOCKING

---

<b>Default value:</b>	ON
<b>Range of values:</b>	OFF, ON

---

Controls whether Heterogeneous Services attempts to optimize performance of data transfer between the Oracle database server and the Heterogeneous Services agent connected to the non-Oracle data store.

The following values are possible:

- `OFF` disables reblocking of fetched data so that data is immediately sent from agent to server
- `ON` enables reblocking, which means that data fetched from the non-Oracle system is buffered in the agent and is not sent to the Oracle database server until the amount of fetched data is equal or higher than `HS_RPC_FETCH_SIZE`. However, any buffered data is returned immediately when a fetch indicates that no more data exists or when the non-Oracle system reports an error.

## HS\_RPC\_FETCH\_SIZE

---

<b>Default value:</b>	4000
<b>Range of values:</b>	Decimal integer (byte count)

---

Tunes internal data buffering to optimize the data transfer rate between the server and the agent process.

Increasing the value can reduce the number of network round trips needed to transfer a given amount of data, but also tends to increase data bandwidth and to reduce response time or **latency** as measured between issuing a query and completion of all fetches for the query. Nevertheless, increasing the fetch size can increase latency for the initial fetch results of a query, because the first fetch results are not transmitted until additional data is available.

After the gateway is installed and configured, you can use the gateway to access non-Oracle database system data, pass non-Oracle database system commands

---

from applications to the non-Oracle database system database, perform distributed queries, and copy data.

## HS\_TIME\_ZONE

---

**Default value for** '[+ | -] hh:mm': Derived from NLS\_TERRITORY

**Range of values for** Any valid datetime format mask  
'[+ | -] hh:mm':

---

Specifies the default local time zone displacement for the current SQL session. The format mask, [+ | -]hh:mm, is specified to indicate the hours and minutes before or after UTC (Coordinated Universal Time—formerly Greenwich Mean Time) For example:

```
HS_TIME_ZONE = [+ | -] hh:mm
```

## IFILE

---

**Default value:** None

**Range of values:** Valid parameter filenames

---

Use IFILE to embed another initialization file within the current initialization file; the value should be an absolute path and should not contain environment variables; the three levels of nesting limit does not apply.

**See Also:** IFILE in *Oracle9i Database Reference*.



---

## Data Type Mapping

Oracle9i maps the ANSI data types through ODBC and OLE DB interfaces to supported Oracle data types. When the results of a query are returned, Oracle9i converts the ODBC or OLE DB data types to Oracle data types.

The tables in this appendix show how Oracle maps ANSI data types through ODBC and OLE DB interfaces to supported Oracle data types when it is retrieving data from a non-Oracle system.

This appendix contains the following tables

- [Mapping ANSI Data Types to Oracle Data Types Through an ODBC Interface](#)
- [Mapping ANSI Data Types to Oracle Data Types Through an OLE DB Interface](#)

## Mapping ANSI Data Types to Oracle Data Types Through an ODBC Interface

**Table 7-4 Mapping ANSI Data Types to Oracle Data Types Through an ODBC Interface**

ANSI	ODBC	Oracle
NUMERIC(19,0)	SQL_BIGINT	NUMBER(19,0)
N/A	SQL_BINARY	RAW
CHAR	SQL_CHAR	CHAR
DATE	SQL_DATE	DATE
DECIMAL(p,s)	SQL_DECIMAL(p,s)	NUMBER(p,s)
DOUBLE PRECISION	SQL_DOUBLE	FLOAT(49)
FLOAT	SQL_FLOAT	FLOAT(49)
INTEGER	SQL_INTEGER	NUMBER(10) <sup>1</sup>
N/A	SQL_LONGVARBINARY	LONG RAW
N/A	SQL_LONGVARCHAR	LONG <sup>2</sup>
REAL	SQL_REAL	FLOAT(23)
SMALLINT	SQL_SMALLINT	NUMBER(5)
TIME	SQL_TIME	DATE
TIMESTAMP	SQL_TIMESTAMP	DATE
NUMERIC(3,0)	SQL_TINYINT	NUMBER(3)
VARCHAR	SQL_VARCHAR	VARCHAR

<sup>1</sup> It's possible under some circumstance for the INTEGER ANSI data type to map to Precision 38, but it usually maps to Precision 10.

<sup>2</sup> If an ANSI SQL implementation defines a large value for the maximum length of VARCHAR data, then it is possible that ANSI VARCHAR will map to SQL\_LONGVARCHAR and Oracle LONG. The same is true for OLE DB DBTYPE\_STRING (long attribute).

---

---

**Note:** This table maps ODBC data types into equivalent ANSI and Oracle data types. In some cases equivalence to ANSI data types is not guaranteed to be exact because the ANSI SQL standard delegates definition of numeric precision and maximum length of character data to individual implementations. This table reflects a probable mapping between ANSI and ODBC data types for a typical implementation of ANSI SQL.

---

---

## Mapping ANSI Data Types to Oracle Data Types Through an OLE DB Interface

*Table 7-5*

ANSI	OLE DB	Oracle
NUMERIC(3,0)	DBTYPE_UI1	NUMBER(3)
NUMERIC(3,0)	DBTYPE_I1	NUMBER(3)
SMALLINT	DBTYPE_UI2	NUMBER(5)
SMALLINT	DBTYPE_I2	NUMBER(5)
NUMERIC(3,0)	DBTYPE_BOOL	NUMBER(5)
INTEGER	DBTYPE_UI4	NUMBER(10)
INTEGER	DBTYPE_I4	NUMBER(10)
NUMERIC(19,0)	DBTYPE_UI8	NUMBER(19,0)
NUMERIC(19,0)	DBTYPE_I8	NUMBER(19,0)
NUMERIC(p,s)	DBTYPE_NUMERIC(p,s)	NUMBER(p,s)
FLOAT	DBTYPE_R4	FLOAT(23)
DOUBLE PRECISION	DBTYPE_R8	FLOAT(49)
N/A	DBTYPE_DECIMAL	FLOAT(49)
VARCHAR	DBTYPE_STR	VARCHAR2
VARCHAR	DBTYPE_WSTR	VARCHAR2
NUMERIC(19,0)	DBTYPE_CY	NUMBER(19,0)
DATE	DBTYPE_DBDATE	DATE
TIME	DBTYPE_DBTIME	DATE
TIMESTAMP	DBTYPE_TIMESTAMP	DATE
N/A	DBTYPE_BYTES	RAW
N/A	DBTYPE_BYTES (long attribute)	LONG RAW
N/A	DBTYPE_STRING (long attribute)	LONG

---

---

# DBMS\_HS\_PASSTHROUGH for Pass-Through SQL

The package, `DBMS_HS_PASSTHROUGH`, contains the procedures and functions for pass-through SQL of Heterogeneous Services. This appendix documents each of them.

This appendix contains these topics:

- `BIND_VARIABLE` procedure
- `BIND_VARIABLE_NCHAR` procedure
- `BIND_VARIABLE_RAW` procedure
- `BIND_OUT_VARIABLE` procedure
- `BIND_OUT_VARIABLE_NCHAR` procedure
- `BIND_OUT_VARIABLE_RAW` procedure
- `BIND_INOUT_VARIABLE` procedure
- `BIND_INOUT_VARIABLE_NCHAR` procedure
- `BIND_INOUT_VARIABLE_RAW` procedure
- `CLOSE_CURSOR` function
- `EXECUTE_IMMEDIATE` function
- `EXECUTE_NON_QUERY` function
- `FETCH_ROW` function
- `GET_VALUE` procedure
- `GET_VALUE_NCHAR` procedure

- 
- `GET_VALUE_RAW` procedure
  - `OPEN_CURSOR` function
  - `PARSE` procedure

## Summary of Subprograms

**Table C-1** DBMS\_HS Package Subprograms

Subprogram	Description
<a href="#">BIND_VARIABLE procedure</a>	Binds an IN variable positionally with a PL/SQL program variable.
<a href="#">BIND_VARIABLE_NCHAR procedure</a>	Binds IN variables of type NVARCHAR2.
<a href="#">BIND_VARIABLE_RAW procedure</a>	Binds IN variables of type RAW.
<a href="#">BIND_OUT_VARIABLE procedure</a>	Binds an OUT variable with a PL/SQL program variable.
<a href="#">BIND_OUT_VARIABLE_NCHAR procedure</a>	Binds an OUT variable of data type NVARCHAR2 with a PL/SQL program variable.
<a href="#">BIND_OUT_VARIABLE_RAW procedure</a>	Binds an OUT variable of data type RAW with a PL/SQL program variable.
<a href="#">BIND_INOUT_VARIABLE procedure</a>	Binds IN OUT bind variables.
<a href="#">BIND_INOUT_VARIABLE_NCHAR procedure</a>	Binds IN OUT bind variables of data type NVARCHAR2
<a href="#">BIND_INOUT_VARIABLE_RAW procedure</a>	Binds IN OUT bind variables of data type RAW
<a href="#">CLOSE_CURSOR function</a>	Closes the cursor and releases associated memory after the SQL statement has been executed at the non-Oracle system
<a href="#">EXECUTE_IMMEDIATE function</a>	Executes a SQL statement immediately
<a href="#">EXECUTE_NON_QUERY function</a>	Executes any SQL statement other than a SELECT statement
<a href="#">FETCH_ROW function</a>	Fetches rows from a result set
<a href="#">GET_VALUE procedure</a>	Retrieves the select list items of SELECT statements after a row has been fetched, and retrieves the OUT bind values after the SQL statement has been executed

**Table C-1** DBMS\_HS Package Subprograms

Subprogram	Description
<a href="#">GET_VALUE_NCHAR procedure</a>	Retrieves the select list items of SELECT statements after a row has been fetched, and retrieves the OUT bind values after the SQL statement has been executed. This procedure operates on the NVARCHAR2 data type
<a href="#">GET_VALUE_RAW procedure</a>	Retrieves the select list items of SELECT statements after a row has been fetched, and retrieves the OUT bind values after the SQL statement has been executed. This procedure operates on the RAW data type
<a href="#">OPEN_CURSOR function</a>	Opens a cursor for executing a pass-through SQL statement at the non-Oracle system
<a href="#">PARSE procedure</a>	Parses a SQL statement at non-Oracle system

## BIND\_VARIABLE procedure

See Also:

- [OPEN\\_CURSOR function](#)
- [PARSE procedure](#)
- [BIND\\_VARIABLE\\_NCHAR procedure](#)
- [BIND\\_VARIABLE\\_RAW procedure](#)

This procedure binds an IN variable positionally with a PL/SQL program variable.

### Syntax.

```
DBMS_HS_PASSTHROUGH.BIND_VARIABLE (
  c      IN      BINARY_INTEGER NOT NULL,
  pos    IN      BINARY_INTEGER NOT NULL,
  val    IN      dt);
```

Where *dt* is one of the following data types:

- DATE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND
- NUMBER



- `TIMESTAMP`
- `TIMESTAMP WITH TIMEZONE`
- `TIMESTAMP WITH LOCAL TIMEZONE`
- `VARCHAR2`

**See Also:**

- [BIND\\_VARIABLE\\_NCHAR procedure](#)
- [BIND\\_VARIABLE\\_RAW procedure](#)

## Parameters

**Table C-2** *BIND\_VARIABLE Procedure Parameters*

Parameter	Description
<code>c</code>	Cursor associated with the pass-through SQL statement. The cursor must be opened and parsed using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> .
<code>pos</code>	Position of the bind variable in the SQL statement. Starts from 1
<code>val</code>	Value that must be passed to the bind variable

## Exceptions

**Table C-3** *BIND\_VARIABLE Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid
ORA-28552	The procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range
ORA-28555	A <code>NULL</code> value was passed for a <code>NOT NULL</code> parameter

## Pragmas

Purity levels defined: `WNDS`, `RNDS`

## **BIND\_VARIABLE\_NCHAR procedure**

This procedure binds IN variables of type NVARCHAR2.

### **Syntax**

```
DBMS_HS_PASSTHROUGH.BIND_VARIABLE_NCHAR (  
    c          IN    BINARY_INTEGER NOT NULL,  
    pos       IN    BINARY_INTEGER NOT NULL,  
    val       IN    NVARCHAR2);
```

### **Parameters**

**Table C-4** BIND\_VARIABLE\_NCHAR Procedure Parameters

<b>Parameter</b>	<b>Description</b>
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed using the routines OPEN_CURSOR and PARSE.
pos	Position of the bind variable in the SQL statement. Starts from 1.
val	Value that must be passed to the bind variable

### **Exceptions**

**Table C-5** BIND\_VARIABLE\_NCHAR Procedure Exceptions

<b>Exception</b>	<b>Description</b>
ORA-28550	The cursor passed is invalid
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range
ORA-28555	A NULL value was passed for a NOT NULL parameter

### **Pragmas**

Purity level defined: WNDS, RNDS

**See Also:**

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure

**BIND\_VARIABLE\_RAW** procedure

This procedure binds IN variables of type RAW.

**Syntax**

```
DBMS_HS_PASSTHROUGH.BIND_VARIABLE_RAW (
    c      IN    BINARY_INTEGER NOT NULL,
    pos    IN    BINARY_INTEGER NOT NULL,
    val    IN    RAW);
```

**Parameters****Table C-6** *BIND\_VARIABLE\_RAW Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> .
pos	Position of the bind variable in the SQL statement. Starts from 1.
val	Value that must be passed to the bind variable

**Exceptions****Table C-7** *BIND\_VARIABLE\_RAW Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range
ORA-28555	A NULL value was passed for a NOT NULL parameter

## Pragmas

Purity level defined: WNDS, RNDS

### See Also:

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [BIND\\_VARIABLE](#) procedure
- [BIND\\_OUT\\_VARIABLE](#) procedure

## BIND\_OUT\_VARIABLE procedure

This procedure binds an OUT variable with a PL/SQL program variable.

### Syntax

```
DBMS_HS_PASSTHROUGH.BIND_OUT_VARIABLE (  
    c          IN    BINARY_INTEGER NOT NULL,  
    pos        IN    BINARY_INTEGER NOT NULL,  
    val        OUT   dtv);
```

Where *dtv* is one of

- DATE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND
- NUMBER
- TIMESTAMP
- TIMESTAMP WITH TIMEZONE
- TIMESTAMP WITH LOCAL TIMEZONE
- VARCHAR2

**See Also:** [BIND\\_INOUT\\_VARIABLE\\_NCHAR procedure](#) for more information about OUT variables of data type RAW.

## Parameters

**Table C-8** *BIND\_OUT\_VARIABLE Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
pos	Position of the bind variable in the SQL statement. Starts from 1.
val	Variable in which the OUT bind variable will store its value. The package will remember only the size of the variable. After the SQL statement is executed, you can use GET_VALUE to retrieve the value of the OUT parameter. The size of the retrieved value should not exceed the size of the parameter that was passed using BIND_OUT_VARIABLE.

## Exceptions

**Table C-9** *BIND\_OUT\_VARIABLE Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Purity level defined: WNDS, RNDS

### See Also:

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR](#) procedure
- [BIND\\_VARIABLE](#) procedure
- [BIND\\_VARIABLE\\_NCHAR](#) procedure
- [GET\\_VALUE](#) procedure

## BIND\_OUT\_VARIABLE\_NCHAR procedure

This procedure binds an OUT variable of data type NVARCHAR2 with a PL/SQL program variable.

### Syntax

```
DEMS_HS_PASSTHROUGH.BIND_OUT_VARIABLE (  
    c      IN    BINARY_INTEGER NOT NULL,  
    pos   IN    BINARY_INTEGER NOT NULL,  
    val   OUT   NVARCHAR2);
```

### Parameters

**Table C-10** *BIND\_OUT\_VARIABLE\_NCHAR Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
pos	Position of the bind variable in the SQL statement. Starts from 1.
val	Variable in which the OUT bind variable will store its value. The package will remember only the size of the variable. After the SQL statement is executed, you can use GET_VALUE to retrieve the value of the OUT parameter. The size of the retrieved value should not exceed the size of the parameter that was passed using BIND_OUT_VARIABLE_RAW.

## Exceptions

**Table C-11** *BIND\_OUT\_VARIABLE\_NCHAR Parameter Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Pragmas defined: WNDS , RNDS

## BIND\_OUT\_VARIABLE\_RAW procedure

This procedure binds an OUT variable of data type RAW with a PL/SQL program variable.

## Syntax

```
DBMS_HS_PASSTHROUGH.BIND_OUT_VARIABLE (
    c          IN    BINARY_INTEGER NOT NULL,
    pos       IN    BINARY_INTEGER NOT NULL,
    val       OUT   RAW);
```

## Parameters

**Table C-12** *BIND\_OUT\_VARIABLE\_RAW Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
pos	Position of the bind variable in the SQL statement. Starts from 1.

**Table C-12** *BIND\_OUT\_VARIABLE\_RAW Procedure Parameters*

Parameter	Description
val	Variable in which the OUT bind variable will store its value. The package will remember only the size of the variable. After the SQL statement is executed, you can use GET_VALUE to retrieve the value of the OUT parameter. The size of the retrieved value should not exceed the size of the parameter that was passed using BIND_OUT_VARIABLE_RAW.

## Exceptions

**Table C-13** *BIND\_OUT\_VARIABLE\_RAW Parameter Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Pragmas defined: WNDS , RNDS

### See Also:

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [BIND\\_OUT\\_VARIABLE](#) procedure
- [BIND\\_VARIABLE](#) procedure
- [BIND\\_VARIABLE\\_NCHAR](#) procedure
- [GET\\_VALUE](#) procedure

## **BIND\_INOUT\_VARIABLE** procedure

This procedure binds IN OUT bind variables.



## Syntax

```
DBMS_HS_PASSTHROUGH.BIND_INOUT_VARIABLE (
    c          IN      BINARY_INTEGER NOT NULL,
    pos       IN      BINARY_INTEGER NOT NULL,
    val       IN OUT  <dt>);
```

Where *dt* is one of

- DATE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND
- NUMBER
- TIMESTAMP
- TIMESTAMP WITH TIMEZONE
- TIMESTAMP WITH LOCAL TIMEZONE
- VARCHAR2

## Parameters

**Table C-14** *BIND\_INOUT\_VARIABLE Procedure Parameters*

Parameter	Description
<i>c</i>	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> respectively.
<i>pos</i>	Position of the bind variable in the SQL statement. Starts from 1.
<i>val</i>	This value will be used for two purposes: <ul style="list-style-type: none"> <li>■ To provide the IN value before the SQL statement is executed</li> <li>■ To determine the size of the OUT value</li> </ul>

## Exceptions

**Table C-15** *BIND\_INOUT\_VARIABLE Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Purity level defined: WNDS, RNDS

### See Also:

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR](#) procedure
- [BIND\\_OUT\\_VARIABLE](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR](#) procedure
- [BIND\\_VARIABLE](#) procedure
- [BIND\\_VARIABLE\\_NCHAR](#) procedure
- [GET\\_VALUE](#) procedure

## BIND\_INOUT\_VARIABLE\_NCHAR procedure

This procedure binds IN OUT bind variables of data type NVARCHAR2.

### Syntax

```
DBMS_HS_PASSTHROUGH.BIND_INOUT_VARIABLE_NCHAR (
    c          IN      BINARY_INTEGER NOT NULL,
    pos       IN      BINARY_INTEGER NOT NULL,
    val       IN OUT NVARCHAR2);
```

## Parameters

**Table C-16** *BIND\_INOUT\_VARIABLE\_NCHAR Procedure Parameters*

Parameter	Description
<code>c</code>	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed' using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> respectively.
<code>pos</code>	Position of the bind variable in the SQL statement. Starts from 1.
<code>val</code>	This value will be used for two purposes: <ul style="list-style-type: none"> <li>■ To provide the <code>IN</code> value before the SQL statement is executed</li> <li>■ To determine the size of the out value</li> </ul>

## Exceptions

**Table C-17** *BIND\_INOUT\_VARIABLE\_NCHAR Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A <code>NULL</code> value was passed for a <code>NOT NULL</code> parameter.

## Pragmas

Pragmas defined: `WNDS`, `RNDS`

**See Also:**

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [BIND\\_INOUT\\_VARIABLE](#) procedure
- [BIND\\_OUT\\_VARIABLE](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR](#) procedure
- [BIND\\_VARIABLE](#) procedure
- [BIND\\_VARIABLE\\_NCHAR](#) procedure
- [GET\\_VALUE](#) procedure

## **BIND\_INOUT\_VARIABLE\_RAW procedure**

This procedure binds IN OUT bind variables of data type RAW.

### **Syntax**

```
DBMS_HS_PASSTHROUGH.BIND_INOUT_VARIABLE_RAW (
    c          IN      BINARY_INTEGER NOT NULL,
    pos       IN      BINARY_INTEGER NOT NULL,
    val       IN OUT RAW);
```

### **Parameters**

**Table C-18** BIND\_INOUT\_VARIABLE\_RAW Procedure Parameters

<b>Parameter</b>	<b>Description</b>
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed using the routines OPEN_CURSOR and PARSE respectively.
pos	Position of the bind variable in the SQL statement. Starts from 1.
val	This value will be used for two purposes: <ul style="list-style-type: none"> <li>■ To provide the IN value before the SQL statement is executed</li> <li>■ To determine the size of the OUT value</li> </ul>

## Exceptions

**Table C-19** *BIND\_INOUT\_VARIABLE\_RAW Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Pragmas defined: WNDS, RNDS

### See Also:

- [OPEN\\_CURSOR function](#)
- [PARSE procedure](#)
- [BIND\\_INOUT\\_VARIABLE procedure](#)
- [BIND\\_OUT\\_VARIABLE procedure](#)
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR procedure](#)
- [BIND\\_VARIABLE procedure](#)
- [BIND\\_VARIABLE\\_NCHAR procedure](#)
- [GET\\_VALUE procedure](#)

## CLOSE\_CURSOR function

This function closes the cursor and releases associated memory after the SQL statement has been executed at the non-Oracle system. If the cursor was not open, the operation is a no operation.

### Syntax

```
DBMS_HS_PASSTHROUGH.CLOSE_CURSOR (
    c    IN    BINARY_INTEGER NOT NULL);
```

## Parameter

**Table C–20** *CLOSE\_CURSOR Procedure Parameters*

Parameter	Description
c	Cursor to be released.

## Exceptions

**Table C–21** *CLOSE\_CURSOR Procedure Exceptions*

Exception	Description
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Purity level defined: WNDS , RNDS

**See Also:** [OPEN\\_CURSOR function](#)

## EXECUTE\_IMMEDIATE function

This function executes a SQL statement immediately. Any valid SQL statement except SELECT can be executed immediately, but the statement must not contain any bind variables. The statement is passed in as a VARCHAR2 in the argument. Internally, the SQL statement is executed using the PASSTHROUGH\_SQL protocol sequence of OPEN\_CURSOR, PARSE, EXECUTE\_NON\_QUERY, CLOSE\_CURSOR.

## Syntax

```
EXECUTE_IMMEDIATE ( s IN VARCHAR2 NOT NULL )
RETURN BINARY_INTEGER);
```

## Parameter Description

**Table C–22** *EXECUTE\_IMMEDIATE Procedure Parameters*

Parameter	Description
s	VARCHAR2 variable with the statement to be executed immediately.

**Returns**

The number of rows affected by the execution of the SQL statement.

**Exceptions**

**Table C-23** *EXECUTE\_IMMEDIATE Procedure Exceptions*

Exception	Description
ORA-28544	Max open cursors.
ORA-28551	SQL statement is invalid.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

**Pragmas**

Purity level defined: NONE

**See Also:**

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [EXECUTE\\_NON\\_QUERY](#) function
- [BIND\\_INOUT\\_VARIABLE](#) procedure

**EXECUTE\_NON\_QUERY function**

This function executes any SQL statement other than a `SELECT` statement. A cursor has to be open and the SQL statement has to be parsed before the SQL statement can be executed.

**Syntax**

```
DBMS_HS_PASSTHROUGH.EXECUTE_NON_QUERY (
  c    IN    BINARY_INTEGER NOT NULL)
RETURN BINARY_INTEGER);
```

## Parameter

**Table C–24** *EXECUTE\_NON\_QUERY Function Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> respectively.

## Returns

The number of rows affected by the SQL statement in the non-Oracle system.

## Exceptions

**Table C–25** *EXECUTE\_NON\_QUERY Function Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	<code>BIND_VARIABLE</code> procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28555	A <code>NULL</code> value was passed for a <code>NOT NULL</code> parameter.

## Pragmas

Purity level defined: NONE

### See Also:

- [OPEN\\_CURSOR function](#)
- [PARSE procedure](#)

## FETCH\_ROW function

This function fetches rows from a result set. The result set is defined with a SQL `SELECT` statement.

Before the rows can be fetched, a cursor has to be opened, and the SQL statement has to be parsed. When there are no more rows to be fetched, the function returns 0. After a 0 return, the `NO_DATA_FOUND` exception occurs when:

- A subsequent `FETCH_ROW` is attempted



- A GET\_VALUE is attempted

## Syntax

```
DBMS_HS_PASSTHROUGH.FETCH_ROW (
    c          IN    BINARY_INTEGER NOT NULL
    [,first IN    BOOLEAN])
RETURN BINARY_INTEGER);
```

## Parameters and Descriptions

**Table C–26** *FETCH\_ROW Function Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
first	Optional parameter to re-execute a SELECT statement. Possible values: <ul style="list-style-type: none"> <li>■ TRUE: re-execute SELECT statement.</li> <li>■ FALSE: fetch the next row, or if executed for the first time execute and fetch rows (default).</li> </ul>

## Returns

The returns the number of rows fetched. The function will return 0 if the last row was already fetched.

## Exceptions

**Table C–27** *FETCH\_ROW Function Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor and parse the SQL statement?
ORA-28555	A NULL value was passed for a NOT NULL parameter.

## Pragmas

Purity level defined: WNDS

**See Also:**

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure

## GET\_VALUE procedure

This procedure has two purposes:

- To retrieve the select list items of `SELECT` statements after a row has been fetched.
- To retrieve the `OUT` bind values after the SQL statement has been executed.

### Syntax

```
DBMS_HS_PASSTHROUGH.GET_VALUE (  
    c      IN      BINARY_INTEGER NOT NULL,  
    pos    IN      BINARY_INTEGER NOT NULL,  
    val    OUT     <dt>);
```

Where *dt* is one of:

- DATE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND
- NUMBER
- TIMESTAMP
- TIMESTAMP WITH TIMEZONE
- TIMESTAMP WITH LOCAL TIMEZONE
- VARCHAR2

For retrieving values of data type `RAW`, see `GET_VALUE_RAW`.

## Parameters

**Table C-28** *GET\_VALUE Procedure Parameters*

Parameter	Description
<code>c</code>	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> respectively.
<code>pos</code>	Position of the bind variable or select list item in the SQL statement. Starts from 1.
<code>val</code>	Variable in which the <code>OUT</code> bind variable or select list item will store its value.

## Exceptions

**Table C-29** *GET\_VALUE Procedure Exceptions*

Exception	Description
ORA-1403	Returns <code>NO_DATA_FOUND</code> exception when executing the <code>GET_VALUE</code> after the last row was fetched (i.e. <code>FETCH_ROW</code> returned 0).
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor, parse and execute (or fetch) the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A <code>NULL</code> value was passed for a <code>NOT NULL</code> parameter.

## Pragmas

Purity level defined: `WNDS`

**See Also:**

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [FETCH\\_ROW](#) function
- [GET\\_VALUE\\_NCHAR](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_RAW](#) procedure

## GET\_VALUE\_NCHAR procedure

This procedure, which operates on NVARCHAR2 data types, has two purposes:

- To retrieve the select list items of `SELECT` statements after a row has been fetched.
- To retrieve the `OUT` bind values after the `SQL` statement has been executed.

### Syntax

```
DBMS_HS_PASSTHROUGH.GET_VALUE_NCHAR (
    c      IN    BINARY_INTEGER NOT NULL,
    pos    IN    BINARY_INTEGER NOT NULL,
    val    OUT   NVARCHAR2);
```

### Parameters

**Table C-30** *GET\_VALUE\_NCHAR Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through <code>SQL</code> statement. Cursor must be opened and parsed, using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> respectively.
pos	Position of the bind variable or select list item in the <code>SQL</code> statement. Starts from 1.
val	Variable in which the <code>OUT</code> bind variable or select list item will store its value.

## Exceptions

**Table C-31** *GET\_VALUE\_NCHAR Procedure Exceptions*

Exception	Description
ORA-1403	Returns <code>NO_DATA_FOUND</code> exception when executing the <code>GET_VALUE</code> after the last row was fetched (i.e. <code>FETCH_ROW</code> returned 0).
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor, parse and execute (or fetch) the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A <code>NULL</code> value was passed for a <code>NOT NULL</code> parameter.

## Pragmas

Purity level defined: WNDS

### See Also:

- [OPEN\\_CURSOR function](#)
- [PARSE procedure](#)
- [FETCH\\_ROW function](#)
- [GET\\_VALUE procedure](#)
- [GET\\_VALUE\\_RAW procedure](#)
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR procedure](#)
- [BIND\\_INOUT\\_VARIABLE\\_RAW procedure](#)

## GET\_VALUE\_RAW procedure

This procedure, which operates on `RAW` data types, has two purposes:

- To retrieve the select list items of `SELECT` statements after a row has been fetched.
- To retrieve the `OUT` bind values after the SQL statement has been executed.

## Syntax

```
DBMS_HS_PASSTHROUGH.GET_VALUE_RAW (
```

```

c      IN      BINARY_INTEGER NOT NULL,
pos    IN      BINARY_INTEGER NOT NULL,
val    OUT     RAW);

```

## Parameters

**Table C-32** *GET\_VALUE\_RAW Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened and parsed, using the routines <code>OPEN_CURSOR</code> and <code>PARSE</code> respectively.
pos	Position of the bind variable or select list item in the SQL statement. Starts from 1.
val	Variable in which the <code>OUT</code> bind variable or select list item will store its value.

## Exceptions

**Table C-33** *GET\_VALUE\_RAW Procedure Exceptions*

Exception	Description
ORA-1403	Returns <code>NO_DATA_FOUND</code> exception when executing the <code>GET_VALUE</code> after the last row was fetched (i.e. <code>FETCH_ROW</code> returned 0).
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not executed in right order. Did you first open the cursor, parse and execute (or fetch) the SQL statement?
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A <code>NULL</code> value was passed for a <code>NOT NULL</code> parameter.

## Pragmas

Purity level defined: `WNDS`

**See Also:**

- [OPEN\\_CURSOR](#) function
- [PARSE](#) procedure
- [FETCH\\_ROW](#) function
- [GET\\_VALUE](#) procedure
- [GET\\_VALUE\\_NCHAR](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_NCHAR](#) procedure
- [BIND\\_INOUT\\_VARIABLE\\_RAW](#) procedure

**OPEN\_CURSOR** function

This function opens a cursor for executing a pass-through SQL statement at the non-Oracle system. This function must be called for any type of SQL statement. The function returns a cursor, which must be used in subsequent calls. This call allocates memory. To deallocate the associated memory, you call the procedure `DBMS_HS_PASSTHROUGH.CLOSE_CURSOR`.

**Syntax**

```
DBMS_HS_PASSTHROUGH.OPEN_CURSOR (
    RETURN    BINARY_INTEGER;
```

**Returns**

The cursor to be used on subsequent procedure and function calls.

**Exceptions**

**Table C-34** *OPEN\_CURSOR* Function Exceptions

Exception	Description
ORA-28554	Maximum number of open cursor has been exceeded. Increase Heterogeneous Services <code>OPEN_CURSORS</code> initialization parameter.

**Pragmas**

Purity level defined: `WNDS`, `RNDS`

**See Also:** [BIND\\_INOUT\\_VARIABLE procedure](#)

## PARSE procedure

This procedure parses a SQL statement at non-Oracle system.

### Syntax

```
DBMS_HS_PASSTHROUGH.GET_VALUE_RAW (  
    c      IN      BINARY_INTEGER NOT NULL,  
    stmt  IN      VARCHAR2      NOT NULL);
```

### Parameters

**Table C–35** *PARSE Procedure Parameters*

Parameter	Description
c	Cursor associated with the pass-through SQL statement. Cursor must be opened using function OPEN_CURSOR.
stmt	Statement to be parsed.

### Exceptions

**Table C–36** *PARSE Procedure Exceptions*

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28551	SQL statement is illegal.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

### Pragmas

Purity level defined: WNDS , RNDS



**See Also:**

- `OPEN_CURSOR` function
- `PARSE` procedure
- `FETCH_ROW` function
- `GET_VALUE` procedure
- `BIND_INOUT_VARIABLE_NCHAR` procedure
- `BIND_INOUT_VARIABLE_NCHAR` procedure



---

# Data Dictionary Translation Support

Data dictionary information is stored in the non-Oracle system as system tables and is accessed through ODBC or OLE DB application programming interfaces (APIs). This appendix documents data dictionary translation support. It explains how to access non-Oracle data dictionaries, lists Heterogeneous Services data dictionary views, describes how to use supported views and tables, and explains data dictionary mapping.

This appendix contains the following topics:

- [Accessing the Non-Oracle Data Dictionary](#)
- [Heterogeneous Services Data Dictionary Views](#)
- [Supported Views and Tables](#)
- [Data Dictionary Mapping](#)

## Accessing the Non-Oracle Data Dictionary

Accessing a non-Oracle data dictionary table or view is identical to accessing a data dictionary in an Oracle database. You issue a `SELECT` statement specifying a database link. The Oracle9i data dictionary view and column names are used to access the non-Oracle data dictionary. Synonyms of supported views are also acceptable.

For example, the following statement queries the data dictionary table `ALL_USERS` to retrieve all users in the non-Oracle system:

```
SQL SELECT * FROM all_users@sid1;
```

When you issue a data dictionary access query, the ODBC or OLE DB agent:

1. Maps the requested table, view, or synonym to one or more ODBC or OLE DB APIs (see "[Data Dictionary Mapping](#)"). The agent translates all data dictionary column names to their corresponding non-Oracle column names within the query.
2. Sends the sequence of APIs to the non-Oracle system.
3. Possibly converts the retrieved non-Oracle data to give it the appearance of the Oracle8i data dictionary table.
4. Passes the data dictionary information from the non-Oracle system table to the Oracle8i.

---

---

**Note:** The values returned when querying the generic connectivity data dictionary may not be the same as the ones returned by the Oracle Enterprise Manager DESCRIBE command.

---

---

## Heterogeneous Services Data Dictionary Views

Heterogeneous Services mapping supports the following list of data dictionary views:

- ALL\_CATALOG
- ALL\_COL\_COMMENTS
- ALL\_COL\_PRIVS
- ALL\_COL\_PRIVS\_MADE
- ALL\_COL\_PRIVS\_RECD
- ALL\_CONSTRAINTS
- ALL\_CONS\_COLUMNS
- ALL\_DB\_LINKS
- ALL\_DEF\_AUDIT\_OPTS
- ALL\_DEPENDENCIES
- ALL\_ERRORS
- ALL\_INDEXES
- ALL\_IND\_COLUMNS
- ALL\_OBJECTS

- ALL\_SEQUENCES
- ALL\_SNAPSHOTS
- ALL\_SOURCE
- ALL\_SYNONYMS
- ALL\_TABLES
- ALL\_TAB\_COLUMNS
- ALL\_TAB\_COMMENTS
- ALL\_TAB\_PRIVS
- ALL\_TAB\_PRIVS\_MADE
- ALL\_TAB\_PRIVS\_RECD
- ALL\_TRIGGERS
- ALL\_USERS
- ALL\_VIEWS
- AUDIT\_ACTIONS
- COLUMN\_PRIVILEGES
- DBA\_CATALOG
- DBA\_COL\_COMMENTS
- DBA\_COL\_PRIVS
- DBA\_OBJECTS
- DBA\_ROLES
- DBA\_ROLE\_PRIVS
- DBA\_SYS\_PRIVS
- DBA\_TABLES
- DBA\_TAB\_COLUMNS
- DBA\_TAB\_COMMENTS
- DBA\_TAB\_PRIVS
- DBA\_USERS
- DICTIONARY

- `DICT_COLUMNS`
- `DUAL`
- `INDEX_STATS`
- `PRODUCT_USER_PROFILE`
- `RESOURCE_COST`
- `ROLE_ROLE_PRIVS`
- `ROLE_SYS_PRIVS`
- `ROLE_TAB_PRIVS`
- `SESSION_PRIVS`
- `SESSION_ROLES`
- `TABLE_PRIVILEGES`
- `USER_AUDIT_OBJECT`
- `USER_AUDIT_SESSION`
- `USER_AUDIT_STATEMENT`
- `USER_AUDIT_TRAIL`
- `USER_CATALOG`
- `USER_CLUSTERS`
- `USER_CLU_COLUMNS`
- `USER_COL_COMMENTS`
- `USER_COL_PRIVS`
- `USER_COL_PRIVS_MADE`
- `USER_COL_PRIVS_RECD`
- `USER_CONSTRAINTS`
- `USER_CONS_COLUMNS`
- `USER_DB_LINKS`
- `USER_DEPENDENCIES`
- `USER_ERRORS`
- `USER_EXTENTS`

- USER\_FREE\_SPACE
- USER\_INDEXES
- USER\_IND\_COLUMNS
- USER\_OBJECTS
- USER\_OBJ\_AUDIT\_OPTS
- USER\_RESOURCE\_LIMITS
- USER\_ROLE\_PRIVS
- USER\_SEGMENTS
- USER\_SEQUENCES
- USER\_SNAPSHOT\_LOGS
- USER\_SOURCE
- USER\_SYNONYMS
- USER\_SYS\_PRIVS
- USER\_TABLES
- USER\_TABLESPACES
- USER\_TAB\_COLUMNS
- USER\_TAB\_COMMENTS
- USER\_TAB\_PRIVS
- USER\_TAB\_PRIVS\_MADE
- USER\_TAB\_PRIVS\_RECD
- USER\_TRIGGERS
- USER\_TS\_QUOTAS
- USER\_USERS
- USER\_VIEWS

## Supported Views and Tables

Generic connectivity supports only these views and tables:

- [ALL\\_CATALOG](#)

- ALL\_COL\_COMMENTS
- ALL\_CONS\_COLUMNS
- ALL\_CONSTRAINTS
- ALL\_IND\_COLUMNS
- ALL\_INDEXES
- ALL\_OBJECTS
- ALL\_TAB\_COLUMNS
- ALL\_TAB\_COMMENTS
- ALL\_TABLES
- ALL\_USERS
- ALL\_VIEWS
- DICTIONARY
- USER\_CATALOG
- USER\_COL\_COMMENTS
- USER\_CONS\_COLUMNS
- USER\_CONSTRAINTS
- USER\_IND\_COLUMNS
- USER\_INDEXES
- USER\_OBJECTS
- USER\_TAB\_COLUMNS
- USER\_TAB\_COMMENTS
- USER\_TABLES
- USER\_USERS
- USER\_VIEWS

If you use an unsupported view, then you receive the Oracle8i message for no rows selected.

If you want to query data dictionary views using `SELECT . . . FROM DBA_*`, first connect as Oracle user `SYSTEM` or `SYS`. Otherwise, you receive the following error message:



ORA-28506: Parse error in data dictionary translation for %s stored in %s

Using generic connectivity, queries of the supported data dictionary tables and views beginning with the characters ALL\_ may return rows from the non-Oracle system when you do not have access privileges for those non-Oracle objects. When querying an Oracle database with the Oracle data dictionary, rows are returned only for those objects you are permitted to access.

## Data Dictionary Mapping

The tables in this section list Oracle data dictionary view names and the equivalent ODBC or OLE DB APIs used.

**Table 7–6 Generic Connectivity Data Dictionary Mapping**

View	ODBC API	OLE DB API
ALL_CATALOG	SQLTables	DBSCHEMA_CATALOGS
ALL_COL_COMMENTS	SQLColumns	DBSCHEMA_COLUMNS
ALL_CONS_COLUMNS	SQLPrimaryKeys, SQLForeignKeys	DBSCHEMA_PRIMARY_KEYS, DBSCHEMA_FOREIGN_KEYS
ALL_CONSTRAINTS	SQLPrimaryKeys, SQLForeignKeys	DBSCHEMA_PRIMARY_KEYS, DBSCHEMA_FOREIGN_KEYS
ALL_IND_COLUMNS	SQLStatistics	DBSCHEMA_STATISTICS
ALL_INDEXES	SQLStatistics	DBSCHEMA_STATISTICS
ALL_OBJECTS	SQLTables, SQLProcedures, SQLStatistics	DBSCHEMA_TABLES, DBSCHEMA_PROCEDURES, DBSCHEMA_STATISTICS
ALL_TAB_COLUMNS	SQLColumns	DBSCHEMA_COLUMNS
ALL_TAB_COMMENTS	SQLTables	DBSCHEMA_TABLES
ALL_TABLES	SQLStatistics	DBSCHEMA_STATISTICS
ALL_USERS	SQLTables	DBSCHEMA_TABLES
ALL_VIEWS	SQLTables	DBSCHEMA_TABLES
DICTIONARY	SQLTables	DBSCHEMA_TABLES
USER_CATALOG	SQLTables	DBSCHEMA_TABLES
USER_COL_COMMENTS	SQLColumns	DBSCHEMA_COLUMNS

**Table 7-6 Generic Connectivity Data Dictionary Mapping**

View	ODBC API	OLE DB API
<a href="#">USER_CONS_COLUMNS</a>	SQLPrimaryKeys, SQLForeignKeys	DBSCHEMA_PRIMARY_KEYS, DBSCHEMA_FOREIGN_KEYS
<a href="#">USER_CONSTRAINTS</a>	SQLPrimaryKeys, SQLForeignKeys	DBSCHEMA_PRIMARY_KEYS, DBSCHEMA_FOREIGN_KEYS
<a href="#">USER_IND_COLUMNS</a>	SQLStatistics	DBSCHEMA_STATISTICS
<a href="#">USER_INDEXES</a>	SQLStatistics	DBSCHEMA_STATISTICS
<a href="#">USER_OBJECTS</a>	SQLTables, SQLProcedures, SQLStatistics	DBSCHEMA_TABLES, DBSCHEMA_PROCEDURES, DBSCHEMA_STATISTICS
<a href="#">USER_TAB_COLUMNS</a>	SQLColumns	DBSCHEMA_COLUMNS
<a href="#">USER_TAB_COMMENTS</a>	SQLTables	DBSCHEMA_TABLES
<a href="#">USER_TABLES</a>	SQLStatistics	DBSCHEMA_STATISTICS
<a href="#">USER_USERS</a>	SQLTables	DBSCHEMA_TABLES
<a href="#">USER_VIEWS</a>	SQLTables	DBSCHEMA_TABLES

## Generic Connectivity Data Dictionary Descriptions

The generic connectivity data dictionary tables and views provide this information:

- Name, data type, and width of each column
- The contents of columns with fixed values

In the descriptions that follow, the values in the Null? column may differ from the Oracle9i data dictionary tables and views. Any default value is shown to the right of an item.

### **ALL\_CATALOG**

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLE_TYPE		VARCHAR2 ( 11 )	"TABLE" or "VIEW" or SYNONYM

**ALL\_COL\_COMMENTS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
COLUMN_NAME	NOT NULL	VARCHAR2 ( 30 )	
COMMENTS		VARCHAR2 ( 4000 )	NULL

**ALL\_CONS\_COLUMNS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
CONSTRAINT_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
COLUMN_NAME		VARCHAR2 ( 4000 )	
POSITION		NUMBER	

**ALL\_CONSTRAINTS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
CONSTRAINT_NAME	NOT NULL	VARCHAR2 ( 30 )	
CONSTRAINT_TYPE		VARCHAR2 ( 1 )	"R" or "P"
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
SEARCH_CONDITION		LONG	NULL
R_OWNER		VARCHAR2 ( 30 )	
R_CONSTRAINT_NAME		VARCHAR2 ( 30 )	
DELETE_RULE		VARCHAR2 ( 9 )	"CASCADE" or "NO ACTION" or "SET NULL"
STATUS		VARCHAR2 ( 8 )	NULL

Name	Null?	Type	Value
DEFERRABLE		VARCHAR2 ( 14 )	NULL
DEFERRED		VARCHAR2 ( 9 )	NULL
VALIDATED		VARCHAR2 ( 13 )	NULL
GENERATED		VARCHAR2 ( 14 )	NULL
BAD		VARCHAR2 ( 3 )	NULL
RELY		VARCHAR2 ( 4 )	NULL
LAST_CHANGE		DATE	NULL

**ALL\_IND\_COLUMNS**

Name	Null?	Type	Value
INDEX_OWNER	NOT NULL	VARCHAR2 ( 30 )	
INDEX_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLE_OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
COLUMN_NAME		VARCHAR2 ( 4000 )	
COLUMN_POSITION	NOT NULL	NUMBER	
COLUMN_LENGTH	NOT NULL	NUMBER	
DESCEND		VARCHAR2 ( 4 )	"DESC" or "ASC"

**ALL\_INDEXES**

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2 ( 30 )	
INDEX_NAME	NOT NULL	VARCHAR2 ( 30 )	
INDEX_TYPE		VARCHAR2 ( 27 )	NULL
TABLE_OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
TABLE_TYPE		CHAR ( 5 )	"TABLE "
UNIQUENESS		VARCHAR2 ( 9 )	"UNIQUE " or "NONUNIQUE "
COMPRESSION		VARCHAR2 ( 8 )	NULL
PREFIX_LENGTH		NUMBER	0
TABLESPACE_NAME		VARCHAR2 ( 30 )	NULL
INI_TRANS		NUMBER	0
MAX_TRANS		NUMBER	0
INITIAL_EXTENT		NUMBER	0
NEXT_EXTENT		NUMBER	0
MIN_EXTENTS		NUMBER	0
MAX_EXTENTS		NUMBER	0
PCT_INCREASE		NUMBER	0
PCT_THRESHOLD		NUMBER	0
INCLUDE_COLUMNS		NUMBER	0
FREELISTS		NUMBER	0
FREELIST_GROUPS		NUMBER	0
PCT_FREE		NUMBER	0
LOGGING		VARCHAR2 ( 3 )	NULL
BLEVEL		NUMBER	0
LEAF_BLOCKS		NUMBER	0
DISTINCT_KEYS		NUMBER	
AVG_LEAF_BLOCKS_PER_KEY		NUMBER	0
AVG_DATA_BLOCKS_PER_KEY		NUMBER	0
CLUSTERING_FACTOR		NUMBER	0
STATUS		VARCHAR2 ( 8 )	NULL
NUM_ROWS		NUMBER	0
SAMPLE_SIZE		NUMBER	0

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
LAST_ANALYZED		DATE	NULL
DEGREE		VARCHAR2 ( 40 )	NULL
INSTANCES		VARCHAR2 ( 40 )	NULL
PARTITIONED		VARCHAR2 ( 3 )	NULL
TEMPORARY		VARCHAR2 ( 1 )	NULL
GENERATED		VARCHAR2 ( 1 )	NULL
SECONDARY		VARCHAR2 ( 1 )	NULL
BUFFER_POOL		VARCHAR2 ( 7 )	NULL
USER_STATS		VARCHAR2 ( 3 )	NULL
DURATION		VARCHAR2 ( 15 )	NULL
PCT_DIRECT_ACCESS		NUMBER	0
ITYP_OWNER		VARCHAR2 ( 30 )	NULL
ITYP_NAME		VARCHAR2 ( 30 )	NULL
PARAMETERS		VARCHAR2 ( 1000 )	NULL
GLOBAL_STATS		VARCHAR2 ( 3 )	NULL
DOMIDX_STATUS		VARCHAR2 ( 12 )	NULL
DOMIDX_OPSTATUS		VARCHAR2 ( 6 )	NULL
FUNCIDX_STATUS		VARCHAR2 ( 8 )	NULL

**ALL\_OBJECTS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
OBJECT_NAME	NOT NULL	VARCHAR2 ( 30 )	
SUBOBJECT_NAME		VARCHAR2 ( 30 )	NULL
OBJECT_ID	NOT NULL	NUMBER	0
DATA_OBJECT_ID		NUMBER	0

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OBJECT_TYPE		VARCHAR2 (18)	"TABLE" or "VIEW" or "SYNONYM" or "INDEX" or "PROCEDURE"
CREATED	NOT NULL	DATE	NULL
LAST_DDL_TIME	NOT NULL	DATE	NULL
TIMESTAMP		VARCHAR2 (19)	NULL
STATUS		VARCHAR2 (7)	NULL
TEMPORARY		VARCHAR2 (1)	NULL
GENERATED		VARCHAR2 (1)	NULL
SECONDARY		VARCHAR2 (1)	NULL

**ALL\_TAB\_COLUMNS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 (30)	
TABLE_NAME	NOT NULL	VARCHAR2 (30)	
COLUMN_NAME	NOT NULL	VARCHAR2 (30)	
DATA_TYPE		VARCHAR2 (106)	
DATA_TYPE_MOD		VARCHAR2 (3)	NULL
DATA_TYPE_OWNER		VARCHAR2 (30)	NULL
DATA_LENGTH	NOT NULL	NUMBER	
DATA_PRECISION		NUMBER	
DATA_SCALE		NUMBER	
NULLABLE		VARCHAR2 (1)	"Y" or "N"
COLUMN_ID	NOT NULL	NUMBER	
DEFAULT_LENGTH		NUMBER	0
DATA_DEFAULT		LONG	NULL
NUM_DISTINCT		NUMBER	0

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
LOW_VALUE		RAW ( 32 )	NULL
HIGH_VALUE		RAW ( 32 )	NULL
DENSITY		NUMBER	0
NUM_NULLS		NUMBER	0
NUM_BUCKETS		NUMBER	0
LAST_ANALYZED		DATE	NULL
SAMPLE_SIZE		NUMBER	0
CHARACTER_SET_NAME		VARCHAR2 ( 44 )	NULL
CHAR_COL_DEC_LENGTH		NUMBER	0
GLOBAL_STATS		VARCHAR2 ( 3 )	NULL
USER_STATS		VARCHAR2 ( 3 )	NULL
AVG_COL_LEN		NUMBER	0

**ALL\_TAB\_COMMENTS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLE_TYPE		VARCHAR2 ( 11 )	"TABLE" or "VIEW"
COMMENTS		VARCHAR2 ( 4000 )	NULL

**ALL\_TABLES**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLESPACE_NAME		VARCHAR2 ( 30 )	NULL
CLUSTER_NAME		VARCHAR2 ( 30 )	NULL



<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
IOT_NAME		VARCHAR2 ( 30 )	NULL
PCT_FREE		NUMBER	0
PCT_USED		NUMBER	0
INI_TRANS		NUMBER	0
MAX_TRANS		NUMBER	0
INITIAL_EXTENT		NUMBER	0
NEXT_EXTENT		NUMBER	0
MIN_EXTENTS		NUMBER	0
MAX_EXTENTS		NUMBER	0
PCT_INCREASE		NUMBER	0
FREELISTS		NUMBER	0
FREELIST_GROUPS		NUMBER	0
LOGGING		VARCHAR2 ( 3 )	NULL
BACKED_UP		VARCHAR2 ( 1 )	NULL
NUM_ROWS		NUMBER	
BLOCKS		NUMBER	
EMPTY_BLOCKS		NUMBER	0
AVG_SPACE		NUMBER	0
CHAIN_CNT		NUMBER	0
AVG_ROW_LEN		NUMBER	0
AVG_SPACE_FREELIST_BLOCKS		NUMBER	0
NUM_FREELIST_BLOCKS		NUMBER	0
DEGREE		VARCHAR2 ( 10 )	NULL
INSTANCES		VARCHAR2 ( 10 )	NULL
CACHE		VARCHAR2 ( 5 )	NULL
TABLE_LOCK		VARCHAR2 ( 8 )	NULL
SAMPLE_SIZE		NUMBER	0
LAST_ANALYZED		DATE	NULL

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
PARTITIONED		VARCHAR2 ( 3 )	NULL
IOT_TYPE		VARCHAR2 ( 12 )	NULL
TEMPORARY		VARHCHAR2 ( 1 )	NULL
SECONDARY		VARCHAR2 ( 1 )	NULL
NESTED		VARCHAR2 ( 3 )	NULL
BUFFER_POOL		VARCHAR2 ( 7 )	NULL
ROW_MOVEMENT		VARCHAR2 ( 8 )	NULL
GLOBAL_STATS		VARCHAR2 ( 3 )	NULL
USER_STATS		VARCHAR2 ( 3 )	NULL
DURATION		VARHCHAR2 ( 15 )	NULL
SKIP_CORRUPT		VARCHAR2 ( 8 )	NULL
MONITORING		VARCHAR2 ( 3 )	NULL

**ALL\_USERS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
USERNAME	NOT NULL	VARCHAR2 ( 30 )	
USER_ID	NOT NULL	NUMBER	0
CREATED	NOT NULL	DATE	NULL

**ALL\_VIEWS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
OWNER	NOT NULL	VARCHAR2 ( 30 )	
VIEW_NAME	NOT NULL	VARCHAR2 ( 30 )	
TEXT_LENGTH		NUMBER	0
TEXT	NOT NULL	LONG	NULL
TYPE_TEXT_LENGTH		NUMBER	0
TYPE_TEXT		VARCHAR2 ( 4000 )	NULL

Name	Null?	Type	Value
OID_TEXT_LENGTH		NUMBER	0
OID_TEXT		VARCHAR2(4000)	NULL
VIEW_TYPE_OWNER		VARCHAR2(30)	NULL
VIEW_TYPE		VARCHAR2(30)	NULL

**DICTIONARY**

Name	Null?	Type	Value
TABLE_NAME		VARCHAR2(30)	
COMMENTS		VARCHAR2(4000)	NULL

**USER\_CATALOG**

Name	Null?	Type	Value
TABLE_NAME	NOT NULL	VARCHAR2(30)	
TABLE_TYPE		VARCHAR2(11)	"TABLE" or, "VIEW" or "SYNONYM"

**USER\_COL\_COMMENTS**

Name	Null?	Type	Value
TABLE_NAME	NOT NULL	VARCHAR2(30)	
COLUMN_NAME	NOT NULL	VARCHAR2(30)	
COMMENTS		VARCHAR2(4000)	NULL

**USER\_CONS\_COLUMNS**

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2(30)	
CONSTRAINT_NAME	NOT NULL	VARCHAR2(30)	

Name	Null?	Type	Value
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
COLUMN_NAME		VARCHAR2 ( 4000 )	
POSITION		NUMBER	

**USER\_CONSTRAINTS**

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2 ( 30 )	
CONSTRAINT_NAME	NOT NULL	VARCHAR2 ( 30 )	
CONSTRAINT_TYPE		VARCHAR2 ( 1 )	R or P
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
SEARCH_CONDITION		LONG	NULL
R_OWNER		VARCHAR2 ( 30 )	
R_CONSTRAINT_NAME		VARCHAR2 ( 30 )	
DELETE_RULE		VARCHAR2 ( 9 )	"CASCADE " or "NOACTION" or "SET NULL"
STATUS		VARCHAR2 ( 8 )	NULL
DEFERRABLE		VARCHAR2 ( 14 )	NULL
DEFERRED		VARCHAR2 ( 9 )	NULL
VALIDATED		VARCHAR2 ( 13 )	NULL
GENERATED		VARCHAR2 ( 14 )	NULL
BAD		VARCHAR2 ( 3 )	NULL
RELY		VARCHAR2 ( 4 )	NULL
LAST_CHANGE		DATE	NULL

**USER\_IND\_COLUMNS**

Name	Null?	Type	Value
INDEX_NAME		VARCHAR2 ( 30 )	

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
TABLE_NAME		VARCHAR2 ( 30 )	
COLUMN_NAME		VARCHAR2 ( 4000 )	
COLUMN_POSITION		NUMBER	
COLUMN_LENGTH		NUMBER	
DESCEND		VARCHAR2 ( 4 )	"DESC " or "ASC "

**USER\_INDEXES**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
INDEX_NAME	NOT NULL	VARCHAR2 ( 30 )	
INDEX_TYPE		VARCHAR2 ( 27 )	NULL
TABLE_OWNER	NOT NULL	VARCHAR2 ( 30 )	
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLE_TYPE		VARCHAR2 ( 11 )	"TABLE "
UNIQUENESS		VARCHAR2 ( 9 )	"UNIQUE " or "NONUNIQUE "
COMPRESSION		VARCHAR2 ( 8 )	NULL
PREFIX_LENGTH		NUMBER	0
TABLESPACE_NAME		VARCHAR2 ( 30 )	NULL
INI_TRANS		NUMBER	0
MAX_TRANS		NUMBER	0
INITIAL_EXTENT		NUMBER	0
NEXT_EXTENT		NUMBER	0
MIN_EXTENTS		NUMBER	0
MAX_EXTENTS		NUMBER	0
PCT_INCREASE		NUMBER	0
PCT_THRESHOLD		NUMBER	0
INCLUDE_COLUMNS		NUMBER	0

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
FREELISTS		NUMBER	0
FREELIST_GROUPS		NUMBER	0
PCT_FREE		NUMBER	0
LOGGING		VARCHAR2 ( 3 )	NULL
BLEVEL		NUMBER	0
LEAF_BLOCKS		NUMBER	0
DISTINCT_KEYS		NUMBER	
AVG_LEAF_BLOCKS_PER_KEY		NUMBER	0
AVG_DATA_BLOCKS_PER_KEY		NUMBER	0
CLUSTERING_FACTOR		NUMBER	0
STATUS		VARCHAR2 ( 8 )	NULL
NUM_ROWS		NUMBER	0
SAMPLE_SIZE		NUMBER	0
LAST_ANALYZED		DATE	NULL
DEGREE		VARCHAR2 ( 40 )	NULL
INSTANCES		VARCHAR2 ( 40 )	NULL
PARTITIONED		VARCHAR2 ( 3 )	NULL
TEMPORARY		VARCHAR2 ( 1 )	NULL
GENERATED		VARCHAR2 ( 1 )	NULL
SECONDARY		VARCHAR2 ( 1 )	NULL
BUFFER_POOL		VARCHAR2 ( 7 )	NULL
USER_STATS		VARCHAR2 ( 3 )	NULL
DURATION		VARHCAR2 ( 15 )	NULL
PCT_DIRECT_ACCESS		NUMBER	0
ITYP_OWNER		VARCHAR2 ( 30 )	NULL
ITYP_NAME		VARCHAR2 ( 30 )	NULL
PARAMETERS		VARCHAR2 ( 1000 )	NULL
GLOBAL_STATS		VARCHAR2 ( 3 )	NULL

Name	Null?	Type	Value
DOMIDX_STATUS		VARCHAR2(12)	NULL
DOMIDX_OPSTATUS		VARCHAR2(6)	NULL
FUNCIDX_STATUS		VARCHAR2(8)	NULL

**USER\_OBJECTS**

Name	Null?	Type	Value
OBJECT_NAME		VARCHAR2(128)	
SUBOBJECT_NAME		VARCHAR2(30)	NULL
OBJECT_ID		NUMBER	0
DATA_OBJECT_ID		NUMBER	0
OBJECT_TYPE		VARCHAR2(18)	"TABLE" or "VIEW" or "SYNONYM" or "INDEX" or "PROCEDURE"
CREATED		DATE	NULL
LAST_DDL_TIME		DATE	NULL
TIMESTAMP		VARCHAR2(19)	NULL
STATUS		VARCHAR2(7)	NULL
TEMPORARY		VARCHAR2(1)	NULL
GENERATED		VARCHAR2(1)	NULL
SECONDARY		VARCHAR2(1)	NULL

**USER\_TAB\_COLUMNS**

Name	Null?	Type	Value
TABLE_NAME	NOT NULL	VARCHAR2(30)	
COLUMN_NAME	NOT NULL	VARCHAR2(30)	
DATA_TYPE		VARCHAR2(106)	

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
DATA_TYPE_MOD		VARCHAR2 ( 3 )	NULL
DATA_TYPE_OWNER		VARCHAR2 ( 30 )	NULL
DATA_LENGTH	NOT NULL	NUMBER	
DATA_PRECISION		NUMBER	
DATA_SCALE		NUMBER	
NULLABLE		VARCHAR2 ( 1 )	"Y" or "N"
COLUMN_ID	NOT NULL	NUMBER	
DEFAULT_LENGTH		NUMBER	NULL
DATA_DEFAULT		LONG	NULL
NUM_DISTINCT		NUMBER	NULL
LOW_VALUE		RAW ( 3 2 )	NULL
HIGH_VALUE		RAW ( 3 2 )	NULL
DENSITY		NUMBER	0
NUM_NULLS		NUMBER	0
NUM_BUCKETS		NUMBER	0
LAST_ANALYZED		DATE	NULL
SAMPLE_SIZE		NUMBER	0
CHARACTER_SET_NAME		VARCHAR2 ( 44 )	NULL
CHAR_COL_DECL_LENGTH		NUMBER	0
GLOBAL_STATS		VARCHAR2 ( 3 )	NULL
USER_STATS		VARCHAR2 ( 3 )	NULL
AVG_COL_LEN		NUMBER	0

**USER\_TAB\_COMMENTS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
TABLE_NAME	NOT NULL	VARCHAR2 ( 30 )	
TABLE_TYPE		VARCHAR2 ( 11 )	"TABLE" or "VIEW"



<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
COMMENTS		VARCHAR2(4000)	NULL

**USER\_TABLES**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
TABLE_NAME	NOT NULL	VARCHAR2(30)	
TABLESPACE_NAME		VARCHAR2(30)	NULL
CLUSTER_NAME		VARCHAR2(30)	NULL
IOT_NAME		VARCHAR2(30)	NULL
PCT_FREE		NUMBER	0
PCT_USED		NUMBER	0
INI_TRANS		NUMBER	0
MAX_TRANS		NUMBER	0
INITIAL_EXTENT		NUMBER	0
NEXT_EXTENT		NUMBER	0
MIN_EXTENTS		NUMBER	0
MAX_EXTENTS		NUMBER	0
PCT_INCREASE		NUMBER	0
FREELISTS		NUMBER	0
FREELIST_GROUPS		NUMBER	0
LOGGING		VARCHAR2(3)	NULL
BACKED_UP		VARCHAR2(1)	NULL
NUM_ROWS		NUMBER	
BLOCKS		NUMBER	
EMPTY_BLOCKS		NUMBER	0
AVG_SPACE		NUMBER	0
CHAIN_CNT		NUMBER	0
AVG_ROW_LEN		NUMBER	0

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
AVG_SPACE_FREELIST_BLOCKS		NUMBER	0
NUM_FREELIST_BLOCKS		NUMBER	0
DEGREE		VARCHAR2(10)	NULL
INSTANCES		VARCHAR2(10)	NULL
CACHE		VARCHAR2(5)	NULL
TABLE_LOCK		VARCHAR2(8)	NULL
SAMPLE_SIZE		NUMBER	0
LAST_ANALYZED		DATE	NULL
PARTITIONED		VARCHAR2(3)	NULL
IOT_TYPE		VARCHAR2(12)	NULL
TEMPORARY		VARHCHAR2(1)	NULL
SECONDARY		VARCHAR2(1)	NULL
NESTED		VARCHAR2(3)	NULL
BUFFER_POOL		VARCHAR2(7)	NULL
ROW_MOVEMENT		VARCHAR2(8)	NULL
GLOBAL_STATS		VARCHAR2(3)	NULL
USER_STATS		VARCHAR2(3)	NULL
DURATION		VARCHAR2(15)	NULL
SKIP_CORRUPT		VARCHAR2(8)	NULL
MONITORING		VARCHAR2(3)	NULL

**USER\_USERS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
USERNAME	NOT NULL	VARCHAR2(30)	
USER_ID	NOT NULL	NUMBER	0
ACCOUNT_STATUS	NOT NULL	VARCHAR2(32)	OPEN
LOCK_DATE		DATE	NULL

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
EXPIRY_DATE		DATE	NULL
DEFAULT_TABLESPACE	NOT NULL	VARCHAR2(30)	NULL
TEMPORARY_TABLESPACE	NOT NULL	VARCHAR2(30)	NULL
CREATED	NOT NULL	DATE	NULL
INITIAL_RSRC_CONSUMER_GROUP		VARCHAR2(30)	NULL
EXTERNAL_NAME		VARCHAR2(4000)	NULL

**USER\_VIEWS**

<b>Name</b>	<b>Null?</b>	<b>Type</b>	<b>Value</b>
VIEW_NAME	NOT NULL	VARCHAR2(30)	
TEXT_LENGTH		NUMBER	0
TEXT		LONG	NULL
TYPE_TEXT_LENGTH		NUMBER	0
TYPE_TEXT		VARCHAR2(4000)	NULL
OID_TEXT_LENGTH		NUMBER	0
OID_TEXT		VARCHAR2(4000)	NULL
VIEW_TYPE_OWNER		VARCHAR2(30)	NULL
VIEW_TYPE		VARCHAR2(30)	NULL



---

---

# Index

## A

---

### agents

- generic connectivity, 2-4
- Heterogeneous Services
  - disabling self-registration, 4-15
  - registering, 4-11, 4-12, 4-13
- specifying initialization parameters for, 4-4

### application development

- Heterogeneous Services
  - controlling array fetches between non-Oracle server and agent, 4-9
  - controlling array fetches between Oracle server and agent, 4-9
  - controlling reblocking of array fetches, 4-9
  - DBMS\_HS\_PASSTHROUGH package, 3-5
  - pass-through SQL, 3-5
  - using bulk fetches, 4-8
  - using OCI for bulk fetches, 4-9

### array fetches, 4-9

- agents, 4-9

## B

---

### bind queries

- executing using pass-through SQL, 3-11
- BIND\_INOUT\_VARIABLE procedure, 3-6, 3-10
- BIND\_OUT\_VARIABLE procedure, 3-6, 3-10
- BIND\_VARIABLE procedure, 3-6

### buffers

- multiple rows, 3-12

### bulk fetches

- optimizing data transfers using, 4-8

## C

---

### CATHO.SQL script

- installing data dictionary for Heterogeneous Services, 4-2

### character sets

- Heterogeneous Services, A-7

### CLOSE\_CURSOR function, 3-6

### commit point site

- commit point strength, A-3

### configuring

- generic connectivity, 7-8
- transparent gateways, 4-2

### Copying data

- COPY command, 4-19
- from Oracle database server to SQL Server, 4-19
- from SQL Server to Oracle database server, 4-21
- INSERT statement, 4-20

### CREATE TABLE command, 4-21

## D

---

### data dictionary

- contents with generic connectivity, D-5
- installing for Heterogeneous Services, 4-2
- mapping for generic connectivity, D-7
- Oracle server name/SQL Server name, D-7
- translation support for generic connectivity, D-1

### data dictionary views

- generic connectivity, D-5
- Heterogeneous Services, 4-22, D-2

### database links

- heterogeneous systems, 4-4

- date formats
  - Heterogeneous Services, A-8, A-9
- DBMS\_HS\_PASSTHROUGH package, 3-5
  - list of functions and procedures, 3-6
- DBMS\_HS\_PASSTHROUGH.EXECUTE\_  
IMMEDIATE, C-18
- describe cache high water mark
  - definition, A-4
- drivers
  - ODBC, 7-13
  - OLE DB (FS), 7-16
  - OLE DB (SQL), 7-15
- dynamic performance views
  - Heterogeneous Services, 4-28
    - determining open sessions, 4-28
    - determining which agents are on host, 4-28

## E

---

- EXECUTE\_IMMEDIATE procedure, 3-6
  - restrictions, 3-7
- EXECUTE\_NON\_QUERY procedure, 3-6

## F

---

- FDS\_CLASS, 4-13
- FDS\_CLASS\_VERSION, 4-13
- FDS\_INST\_NAME, 4-14
- FETCH\_ROW procedure, 3-7
  - executing queries using pass-through SQL, 3-11
- fetches
  - bulk, 4-8
  - optimizing round-trips, 3-12

## G

---

- Gateway
  - how it works, 2-8
  - remote data access, 1-5
  - two-phase commit, 1-6
- generic connectivity
  - architecture, 7-3
    - Oracle and non-Oracle on same machine, 7-4
    - Oracle and non-Oracle on separate machines, 7-3

- configuration, 7-8
- creating initialization file, 7-8
- data dictionary
  - translation support, D-1
- definition, 7-2
- DELETE statement, 7-7
- editing initialization file, 7-8
- error tracing, A-6
- Heterogeneous Services, 2-4
- INSERT statement, 7-7
- non-Oracle data dictionary access, D-1
- ODBC connectivity requirements, 7-13
- OLE DB (FS) connectivity requirements, 7-16
- OLE DB (SQL) connectivity requirements, 7-15
- restrictions, 7-6
- setting parameters for ODBC source, 7-10
  - UNIX, 7-11
  - Windows NT, 7-10
- setting parameters for OLE DB source, 7-12
- SQL execution, 7-6
- supported functions, 7-7
- supported SQL syntax, 7-7
- types of agents, 7-2
- UPDATE statement, 7-7
- GET\_VALUE procedure, 3-7, 3-10

## H

---

- heterogeneous distributed systems
  - accessing, 4-2
- Heterogeneous Services
  - agent registration, 4-11
    - avoiding configuration mismatches, 4-12
    - disabling, 4-15
    - enabling, 4-11
  - agents
    - self-registration, 4-13
  - application development
    - controlling array fetches between non-Oracle server and agent, 4-9
    - controlling array fetches between Oracle server and agent, 4-9
    - controlling reblocking of array fetches, 4-9
  - DBMS\_HS\_PASSTHROUGH package, 3-5
  - pass-through SQL, 3-5

- using bulk fetches, 4-8
  - using OCI for bulk fetches, 4-9
- creating database links, 4-4
- data dictionary views, 4-22, D-2
  - types, 4-22
  - understanding sources, 4-23
  - using general views, 4-24
  - using SQL service views, 4-26
  - using transaction service views, 4-25
- defining maximum number of open cursors, A-10
- dynamic performance views, 4-28
  - V\$HS\_AGENT view, 4-28
  - V\$HS\_SESSION view, 4-28
- generic connectivity
  - architecture, 7-3
  - creating initialization file, 7-8
  - definition, 7-2
  - editing initialization file, 7-8
  - non-Oracle data dictionary access, D-1
  - ODBC connectivity requirements, 7-13
  - OLE DB (FS) connectivity requirements, 7-16
  - OLE DB (SQL) connectivity requirements, 7-15
  - restrictions, 7-6
  - setting parameters for ODBC source, 7-10
  - setting parameters for OLE DB source, 7-12
  - SQL execution, 7-6
  - supported functions, 7-7
  - supported SQL syntax, 7-7
  - supported tables, D-5
  - types of agents, 7-2
- installing data dictionary, 4-2
- optimizing data transfer, A-11
- setting global name, A-4
- setting up access using transparent gateway, 4-2
- setting up environment, 4-2
- specifying cache high water mark, A-4
- specifying cache size, A-10
- specifying commit point strength, A-3
- specifying domain, A-3
- specifying instance identifier, A-3
- SQL service, 2-5
- testing connections, 4-4
  - transaction service, 2-4
  - tuning internal data buffering, A-11
- HS\_AUTOREGISTER initialization parameter
  - using to enable agent self-registration, 4-14
- HS\_BASE\_CAPS view, 4-23
- HS\_BASE\_DD view, 4-23
- HS\_CLASS\_CAPS view, 4-23
- HS\_CLASS\_DD view, 4-23
- HS\_CLASS\_INIT view, 4-23
- HS\_COMMIT\_POINT\_STRENGTH initialization parameter, A-3
- HS\_DB\_DOMAIN initialization parameter, A-3
- HS\_DB\_INTERNAL\_NAME initialization parameter, A-3
- HS\_DB\_NAME initialization parameter, A-4
- HS\_DESCRIBE\_CACHE\_HWM initialization parameter, A-4
- HS\_FDS\_CLASS view, 4-23
- HS\_FDS\_CONNECT\_INFO initialization parameter, A-4
  - specifying connection information, 7-9
- HS\_FDS\_FETCH\_ROWS initialization parameter, 4-9
- HS\_FDS\_INST view, 4-23
- HS\_FDS\_SHAREABLE\_NAME initialization parameter, A-6
- HS\_FDS\_TRACE initialization parameter, A-6
- HS\_FDS\_TRACE\_LEVEL initialization parameter
  - enabling agent tracing, 7-9
- HS\_LANGUAGE initialization parameter, A-6
- HS-NLS\_DATE\_FORMAT initialization parameter, A-8
- HS-NLS\_DATE\_LANGUAGE initialization parameter, A-8
- HS-NLS\_NCHAR initialization parameter, A-9
- HS\_OPEN\_CURSORS initialization parameter, A-10
- HS\_ROWID\_CACHE\_SIZE initialization parameter, A-10
- HS\_RPC\_FETCH\_REBLOCKING initialization parameter, 4-10, A-11
- HS\_RPC\_FETCH\_SIZE initialization parameter, 4-9, A-11

## I

---

IFILE, A-12

## L

---

listeners, 4-2

## M

---

multiple rows  
buffering, 3-12

## N

---

National Language Support (NLS)  
Heterogeneous Services, A-6  
character set of non-Oracle source, A-9  
date format, A-8  
languages in character date values, A-8

## O

---

OCI  
optimizing data transfers using, 4-9  
ODBC agents  
connectivity requirements, 7-13  
functions, 7-13  
ODBC connectivity  
data dictionary mapping, D-7  
ODBC driver, 7-13  
requirements, 7-13  
specifying connection information  
UNIX, A-5  
Windows NT, A-5  
specifying path to library, A-6  
OLE DB (FS) drivers, 7-16  
OLE DB (SQL) drivers, 7-15  
OLE DB agents  
connectivity requirements, 7-15, 7-16  
OLE DB connectivity  
data dictionary mapping, D-7  
setting connection information, A-5  
OLE DB drivers  
data provider requirements, 7-16  
initialization properties, 7-18

rowset properties, 7-18  
OPEN\_CURSOR procedure, 3-6  
operating system dependencies, C-1  
Oracle database server  
SQL construct processing, 4-16  
Oracle Net Services listener, 2-3, 4-2  
Oracle precompiler  
optimizing data transfers using, 4-9  
OUT bind variables, 3-10

## P

---

PARSE procedure, 3-6  
pass-through SQL  
avoiding SQL interpretation, 3-5  
executing statements, 3-6  
non-queries, 3-7  
queries, 3-11  
with bind variables, 3-8  
with IN bind variables, 3-9  
with IN OUT bind variables, 3-10  
with OUT bind variables, 3-10  
implications of using, 3-6  
overview, 3-5  
restrictions, 3-6

## Q

---

queries  
pass-through SQL, 3-11

## R

---

reblocking, 4-9  
rows  
buffering multiple, 3-12

## S

---

SELECT statement  
accessing non-Oracle system, D-1  
service names  
specifying in database links, 4-4  
SQL capabilities  
data dictionary tables, 4-26



## SQL service

- data dictionary views, 2-8, 4-22

- Heterogeneous Services, 2-5

- views

  - Heterogeneous Services, 4-26

- Synonyms, 4-17

## T

---

### transaction service

- Heterogeneous Services, 2-4

- views

  - Heterogeneous Services, 4-25

### transparent gateways

- accessing Heterogeneous Services agents, 4-2

- creating database links, 4-4

- installing Heterogeneous Services data

  - dictionary, 4-2

- testing connections, 4-4

- Two-phase commit, 1-6

## U

---

### unsupported functions

- generic connectivity, 7-7

## V

---

### V\$HS\_AGENT view

- determining which agents are on host, 4-28

### V\$HS\_PARAMETER view

- listing HS parameters, 4-29

### V\$HS\_SESSION view

- determining open sessions, 4-29

### variables

- bind, 3-7

