

Oracle® Database

Heterogeneous Connectivity Administrator's Guide

10g Release 1 (10.1)

Part No. B10764-01

December 2003

ORACLE

Oracle Database Heterogeneous Connectivity Administrator's Guide, 10g Release 1 (10.1)

Part No. B10764-01

Copyright © 2001, 2003 Oracle Corporation. All rights reserved.

Primary Author: Cathy Baird

Contributors: Hermann Baer, Jacco Draaijer, Vira Goorah, Govind Lakkoju, Kishan Peyetti

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

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

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

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

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

Oracle is a registered trademark, and Oracle Store, Oracle9i, Oracle8i, PL/SQL, and SQL*Plus are trademarks or registered trademarks of Oracle Corporation. Other names may be trademarks of their respective owners.

Contents

Send Us Your Comments	ix
Preface.....	xi
Audience	xii
Organization.....	xii
Related Documentation	xiv
Conventions.....	xiv
Documentation Accessibility	xix
What's New in Heterogeneous Connectivity?	xxi
HS_CALL_NAME Initialization Parameter.....	xxi
1 Introduction	
The Information Integration Challenge.....	1-2
How Oracle Addresses Synchronous Information Integration.....	1-2
Benefits of Oracle's Solution for Synchronous Information Integration	1-4
Remote Data Can Be Accessed Transparently	1-5
There is No Unnecessary Data Duplication.....	1-5
SQL Statements Can Query Several Different Databases.....	1-5
Oracle's Application Development and End User Tools Can Be Used	1-6
Users Can Talk to a Remote Database in its Own Language.....	1-6

2 The Role of the Heterogeneous Services Component

Heterogeneous Connectivity Process Architecture	2-2
Heterogeneous Services Agents	2-2
Types of Heterogeneous Services Agents	2-3
Oracle Transparent Gateways.....	2-3
Generic Connectivity Agents	2-4
Heterogeneous Services Components	2-4
Transaction Service.....	2-4
SQL Service.....	2-5
Configuring Heterogeneous Services	2-5
Data Dictionary Translations	2-6
Heterogeneous Services Initialization Parameters	2-6
Capabilities	2-6
The Heterogeneous Services Data Dictionary	2-7
Classes and Instances	2-7
Data Dictionary Views	2-8
Gateway Process Flow	2-9

3 Features of Oracle Transparent Gateways and Generic Connectivity

SQL and PL/SQL Support	3-2
Heterogeneous Replication	3-2
Pass-Through SQL	3-5
Using the DBMS_HS_PASSTHROUGH Package	3-5
Considering the Implications of Using Pass-Through SQL	3-6
Executing Pass-Through SQL Statements	3-6
Result Set Support	3-13
Result Set Support In Non-Oracle Systems.....	3-14
Heterogeneous Services Support for Result Sets	3-15
Data Dictionary Translations	3-16
Datetime Datatypes	3-18
Two-Phase Commit Protocol	3-18
Piecewise Long	3-19
SQL*Plus DESCRIBE Command	3-19
Constraints on SQL in a Distributed Environment	3-20
Resolving Remote and Heterogeneous References	3-20

Resolving Important Restrictions.....	3-20
Updates, Inserts, and Deletes.....	3-25
Using Oracle's Optimizer with Heterogeneous Services.....	3-26
Example: Using Index and Table Statistics.....	3-26
Example: Remote Join Optimization.....	3-28
Optimizer Restrictions for Non-Oracle Access.....	3-30

4 Using Heterogeneous Services Agents

Setting Up Access to Non-Oracle Systems.....	4-2
Step 1: Configure Oracle Net Services to Access Heterogeneous Services Agents.....	4-2
Step 2: Create the Database Link to the Non-Oracle System.....	4-3
Step 3: Test the Connection.....	4-4
Setting Initialization Parameters.....	4-5
Name and Location of Heterogeneous Services Initialization Parameter File.....	4-5
Syntax for Initialization Parameter Settings.....	4-5
Gateway Initialization Parameters.....	4-6
Optimizing Data Transfers Using Bulk Fetch.....	4-7
Using OCI, an Oracle Precompiler, or Another Tool for Array Fetches.....	4-8
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
Registering Agents.....	4-10
Enabling Agent Self-Registration.....	4-10
Disabling Agent Self-Registration.....	4-14
Oracle Database Server SQL Construct Processing.....	4-14
Executing User-Defined Functions on a Non-Oracle Database.....	4-15
Using Synonyms.....	4-17
Copying Data from the Oracle Database Server to the Non-Oracle Database System.....	4-18
Copying Data from the Non-Oracle Database System to the Oracle Database Server.....	4-19
Heterogeneous Services Data Dictionary Views.....	4-20
Understanding the Types of Views.....	4-20
Understanding the Sources of Data Dictionary Information.....	4-21
Using the General Views.....	4-22
Using the Transaction Service Views.....	4-23
Using the SQL Service Views.....	4-24

Using the Heterogeneous Services Dynamic Performance Views	4-25
Determining Which Agents Are Running on a Host	4-26
Determining the Open Heterogeneous Services Sessions	4-26
Determining the Heterogeneous Services Parameters	4-27
5 Multithreaded Agents	
Why Use Multithreaded Agents?	5-2
The Challenge of Dedicated Agent Architecture	5-2
The Advantage of Multithreading	5-2
Multithreaded Agent Architecture	5-3
The Monitor Thread	5-5
Dispatcher Threads	5-5
Task Threads.....	5-6
Administering Multithreaded Agents	5-6
Agent Control Utility (agtctl) Commands	5-7
Using Single-Line Command Mode.....	5-8
Using Shell Mode Commands	5-11
Configuration Parameters for Multithreaded Agent Control.....	5-12
6 Performance Tips	
Optimizing Heterogeneous Distributed SQL Statements	6-2
Optimizing Performance of Distributed Queries	6-2
7 Generic Connectivity	
What Is Generic Connectivity?	7-2
Types of Agents.....	7-2
Generic Connectivity Architecture.....	7-2
SQL Execution	7-5
Datatype Mapping.....	7-5
Generic Connectivity Restrictions	7-5
Supported Oracle SQL Statements and Functions	7-6
Configuring Generic Connectivity Agents	7-7
Creating the Initialization File	7-7
Editing the Initialization File.....	7-7

Setting Initialization Parameters for an ODBC-based Data Source	7-9
Setting Initialization Parameters for an OLE DB-based Data Source	7-11
ODBC Connectivity Requirements	7-12
OLE DB (SQL) Connectivity Requirements	7-14
OLE DB (FS) Connectivity Requirements	7-15
OLE DB Interfaces for Data Providers to Expose	7-15
Data Source Properties.....	7-17

A Heterogeneous Services Initialization Parameters

HS_CALL_NAME	A-3
HS_COMMIT_POINT_STRENGTH	A-3
HS_DB_DOMAIN	A-4
HS_DB_INTERNAL_NAME	A-4
HS_DB_NAME	A-5
HS_DESCRIBE_CACHE_HWM	A-5
HS_FDS_CONNECT_INFO	A-5
ODBC-Based Data Source on Windows.....	A-6
ODBC-Based Data Source on UNIX	A-6
OLE DB-Based Data Source (Windows NT Only).....	A-6
HS_FDS_DEFAULT_SCHEMA_NAME	A-6
HS_FDS_SHAREABLE_NAME	A-7
HS_FDS_TRACE_LEVEL	A-7
HS_LANGUAGE	A-7
Character Sets.....	A-8
Language.....	A-8
Territory	A-8
HS_LONG_PIECE_TRANSFER_SIZE	A-9
HS_NLS_DATE_FORMAT	A-9
HS_NLS_DATE_LANGUAGE	A-10
HS_NLS_NCHAR	A-10
HS_NLS_NUMERIC_CHARACTER	A-10
HS_NLS_TIMESTAMP_FORMAT	A-11
HS_NLS_TIMESTAMP_TZ_FORMAT	A-11
HS_OPEN_CURSORS	A-12
HS_ROWID_CACHE_SIZE	A-12

HS_RPC_FETCH_REBLOCKING	A-12
HS_RPC_FETCH_SIZE	A-13
HS_TIME_ZONE	A-14
IFILE	A-14

B Datatype Mapping

Mapping ANSI Datatypes to Oracle Datatypes Through an ODBC Interface	B-2
Mapping ANSI Datatypes to Oracle Datatypes Through an OLE DB Interface	B-3

C Data Dictionary Translation Support

Accessing the Non-Oracle Data Dictionary	C-2
Heterogeneous Services Data Dictionary Views	C-2
Views and Tables Supported by Generic Connectivity	C-5
Data Dictionary Mapping	C-5
Generic Connectivity Data Dictionary Descriptions	C-6

Index

Send Us Your Comments

Oracle Database Heterogeneous Connectivity Administrator's Guide, 10g Release 1 (10.1)

Part No. B10764-01

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
- 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.

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

Preface

This manual describes Oracle's approach for information integration in a heterogeneous environment. Specifically, it describes Oracle Transparent Gateways and Generic Connectivity and is meant to be an administrator's guide for these Oracle products.

This preface contains these topics:

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

Audience

Oracle Database Heterogeneous Connectivity Administrator's Guide is intended for the following users:

- Database administrators who want to administer distributed database systems that involve Oracle to non-Oracle database links
- Application developers who want to make use of the heterogeneous connectivity functionality in the Oracle database server
- Readers who want a high-level understanding of Oracle's architecture for heterogeneous connectivity and how it works.

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

- Relational database concepts and basic database or applications administration as described in the following books:
 - *Oracle Database Concepts*
 - *Oracle Database Administrator's Guide*
 - *Oracle Database Application Developer's Guide - Fundamentals*
- The operating system environment under which database administrators are running Oracle.

Organization

This document contains:

Chapter 1, "Introduction"

This chapter describes the challenges of operating in a heterogeneous environment. Oracle recognizes these challenges and offers both synchronous and asynchronous solutions that enable companies to easily operate in such an environment. The two synchronous solutions, Oracle Transparent Gateways and Generic Connectivity, are discussed in this book.

Chapter 2, "The Role of the Heterogeneous Services Component"

Oracle's synchronous solutions for operating in a heterogeneous environment are Oracle Transparent Gateways and Generic Connectivity. The common component of the Oracle database server for supporting these solutions is Heterogeneous Services. This chapter describes the architecture and functionality of the

Heterogeneous Services component and its interaction with Oracle Transparent Gateways and Generic Connectivity.

Chapter 3, "Features of Oracle Transparent Gateways and Generic Connectivity"

This chapter describes the major features provided by Oracle Transparent Gateways and Generic Connectivity.

Chapter 4, "Using Heterogeneous Services Agents"

This chapter explains how to use Heterogeneous Services agents.

Chapter 5, "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 provides instructions on how to set them.

Appendix B, "Datatype 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, "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:

- *Oracle Database Concepts*
- *Oracle Database Administrator's Guide*
- *Oracle Database Application Developer's Guide - Fundamentals*
- *Oracle Database New Features*

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 *Oracle Database Sample Schemas* for information on how these schemas were created and how you can use them yourself.

Printed documentation is available for sale in the Oracle Store at

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

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

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

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

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

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 for Windows Operating Systems](#)

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
Bold	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 index-organized table .
<i>Italics</i>	Italic typeface indicates book titles or emphasis.	<i>Oracle Database Concepts</i> Ensure that the recovery catalog and target database do <i>not</i> reside on the same disk.
UPPERCASE monospace (fixed-width) font	Uppercase monospace typeface indicates elements supplied by the system. Such elements include parameters, privileges, datatypes, RMAN keywords, SQL keywords, SQL*Plus or utility commands, packages and methods, as well as system-supplied column names, database objects and structures, usernames, and roles.	You can specify this clause only for a NUMBER column. You can back up the database by using the BACKUP command. Query the TABLE_NAME column in the USER_TABLES data dictionary view. Use the DBMS_STATS.GENERATE_STATS procedure.
lowercase monospace (fixed-width) font	Lowercase monospace typeface indicates executables, filenames, directory names, and sample user-supplied elements. Such elements include computer and database names, net service names, and connect identifiers, as well as user-supplied database objects and structures, column names, packages and classes, usernames and roles, program units, and parameter values. Note: Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	Enter sqlplus to open SQL*Plus. The password is specified in the orapwd file. Back up the datafiles and control files in the /disk1/oracle/dbs directory. The department_id, department_name, and location_id columns are in the hr.departments table. Set the QUERY_REWRITE_ENABLED initialization parameter to true. Connect as oe user. The JRepUtil class implements these methods.
lowercase italic monospace (fixed-width) font	Lowercase italic monospace font represents placeholders or variables.	You can specify the <i>parallel_clause</i> . Run <i>Uold_release</i> .SQL where <i>old_release</i> refers to the release you installed prior to upgrading.

Conventions in Code Examples

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

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

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

Convention	Meaning	Example
[]	Brackets enclose one or more optional items. Do not enter the brackets.	DECIMAL (<i>digits</i> [, <i>precision</i>])
{ }	Braces enclose two or more items, one of which is required. Do not enter the braces.	{ENABLE DISABLE}
	A vertical bar represents a choice of two or more options within brackets or braces. Enter one of the options. Do not enter the vertical bar.	{ENABLE DISABLE} [COMPRESS NOCOMPRESS]
...	Horizontal ellipsis points indicate either: <ul style="list-style-type: none"> That we have omitted parts of the code that are not directly related to the example That you can repeat a portion of the code 	CREATE TABLE ... AS <i>subquery</i> ; SELECT <i>col1</i> , <i>col2</i> , ... , <i>coln</i> FROM employees;
. . . .	Vertical ellipsis points indicate that we have omitted several lines of code not directly related to the example.	SQL> SELECT NAME FROM V\$DATAFILE; NAME ----- /fs1/dbs/tbs_01.dbf /fs1/dbs/tbs_02.dbf
Other notation	You must enter symbols other than brackets, braces, vertical bars, and ellipsis points as shown.	acctbal NUMBER(11,2); acct CONSTANT NUMBER(4) := 3;
<i>Italics</i>	Italicized text indicates placeholders or variables for which you must supply particular values.	CONNECT SYSTEM/ <i>system_password</i> DB_NAME = <i>database_name</i>

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

Conventions for Windows Operating Systems

The following table describes conventions for Windows operating systems and provides examples of their use.

Convention	Meaning	Example
Choose Start >	How to start a program.	To start the Database Configuration Assistant, choose Start > Programs > Oracle - <i>HOME_NAME</i> > Configuration and Migration Tools > Database Configuration Assistant.
File and directory names	File and directory names are not case sensitive. The following special characters are not allowed: left angle bracket (<), right angle bracket (>), colon (:), double quotation marks ("), slash (/), pipe (), and dash (-). The special character backslash (\) is treated as an element separator, even when it appears in quotes. If the file name begins with \\, then Windows assumes it uses the Universal Naming Convention.	<pre>c:\winnt"\system32 is the same as C:\WINNT\SYSTEM32</pre>

Convention	Meaning	Example
<code>C:\></code>	Represents the Windows command prompt of the current hard disk drive. The escape character in a command prompt is the caret (^). Your prompt reflects the subdirectory in which you are working. Referred to as the <i>command prompt</i> in this manual.	<code>C:\oracle\oradata></code>
Special characters	The backslash (\) special character is sometimes required as an escape character for the double quotation mark (") special character at the Windows command prompt. Parentheses and the single quotation mark (') do not require an escape character. Refer to your Windows operating system documentation for more information on escape and special characters.	<code>C:\>exp scott/tiger TABLES=emp QUERY="\ "WHERE job='SALESMAN' and sal<1600\ C:\>imp SYSTEM/password FROM USER=scott TABLES=(emp, dept)</code>
<i>HOME_NAME</i>	Represents the Oracle home name. The home name can be up to 16 alphanumeric characters. The only special character allowed in the home name is the underscore.	<code>C:\> net start OracleHOME_NAME\TNSListener</code>

Convention	Meaning	Example
<i>ORACLE_HOME</i> and <i>ORACLE_BASE</i>	<p>In releases prior to Oracle8i release 8.1.3, when you installed Oracle components, all subdirectories were located under a top level <i>ORACLE_HOME</i> directory. For Windows NT, the default location was <code>C:\orant</code> for Windows NT.</p> <p>This release complies with Optimal Flexible Architecture (OFA) guidelines. All subdirectories are not under a top level <i>ORACLE_HOME</i> directory. There is a top level directory called <i>ORACLE_BASE</i> that by default is <code>C:\oracle</code>. If you install the latest Oracle release on a computer with no other Oracle software installed, then the default setting for the first Oracle home directory is <code>C:\oracle\orann</code>, where <i>nn</i> is the latest release number. The Oracle home directory is located directly under <i>ORACLE_BASE</i>.</p> <p>All directory path examples in this guide follow OFA conventions.</p> <p>Refer to <i>Oracle Database Platform Guide for Windows</i> for additional information about OFA compliances and for information about installing Oracle products in non-OFA compliant directories.</p>	Go to the <i>ORACLE_BASE\ORACLE_HOME\rdms\admin</i> directory.

Documentation Accessibility

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

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

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

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

What's New in Heterogeneous Connectivity?

This section describes new features in this release and provides pointers to additional information.

HS_CALL_NAME Initialization Parameter

Use the `HS_CALL_NAME` initialization parameter to specify remote functions that can be referenced by SQL statements.

See Also: ["Executing User-Defined Functions on a Non-Oracle Database"](#) on page 4-15 and ["HS_CALL_NAME"](#) on page A-3

Introduction

This chapter describes the challenges of operating in a heterogeneous environment. Oracle recognizes these challenges and offers both synchronous and asynchronous solutions that enable companies to easily operate in such an environment. The two synchronous solutions, Oracle Transparent Gateways and Generic Connectivity, are discussed in this book.

This chapter contains these topics:

- [The Information Integration Challenge](#)
- [How Oracle Addresses Synchronous Information Integration](#)
- [Benefits of Oracle's Solution for Synchronous Information Integration](#)

See Also: For information about a specific Oracle Transparent Gateway, please consult the Oracle documentation for that specific gateway

The Information Integration Challenge

Information integration 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. This data is just bits and bytes on a file system - and only a database can turn the bits and bytes of data into business information. Integration and consolidation of all business information would allow an organization to easily and quickly take advantage of the synergies inherent in business information.

Consolidation of all data into one database system is often difficult. This is in large part because 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 interoperability becomes possible. The database systems can have different access interfaces, different datatypes, 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 datatypes.

However, there are significant differences as well. SQL dialects can be different, as can transaction semantics. There can be some datatypes 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 for synchronously accessing information from multiple sources.

How Oracle Addresses Synchronous Information Integration

If a client program needs to access or modify data at several Oracle databases, it can open connections to each of them. This approach, however, has several drawbacks. Among them, are the following. If you want to join data from the databases, then the client must contain logic that does that. If data integrity must be guaranteed, then the client will need to contain transaction coordination logic.

Oracle provides another approach called distributed processing, where the client connects to one Oracle database and shifts the burden of joining data and

transaction coordination to that database. The database that the client program connects to is called the local database. Any database other than this one is a remote database. The client program can access objects at any of the remote databases using database links. The Oracle query processor takes care of the joins and its transaction engine takes 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 previously described scenario. From the client's point of view, the remote non-Oracle system functions like a remote Oracle system. It will appear 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 component in the Oracle server called Heterogeneous Services.

The work done by the Heterogeneous Services component is, for the most part, completely transparent to the end user. With only a few exceptions (these are noted in later chapters), you are not required to do anything different to access a non-Oracle system than is required for accessing an Oracle system. The Heterogeneous Services component is used as the foundation for implementing Oracle's access to non-Oracle databases.

The following are two methods that Oracle uses for solving the challenges of information sharing and integration in a heterogeneous environment. Because they are both based on a foundation that is integrated into the database, they can exploit all of the features of the database.

- Oracle Transparent Gateways

An Oracle Transparent Gateway works in conjunction with the Heterogeneous Services component of the Oracle Database server 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 operating on a Sun Solaris platform.

Using an Oracle Transparent Gateway, you can access data anywhere in a distributed database system without being required to know either the location of the data, or how it is stored.

- Generic Connectivity

Oracle provides a set of agents, containing only generic code, that interface with the Heterogeneous Services component and comprise Generic Connectivity. These agents require drivers to provide access to the non-Oracle systems.

Oracle provides Generic Connectivity agents for ODBC and OLE DB that enable you to use ODBC and OLE DB drivers to access non-Oracle databases.

Note: The ODBC and OLE DB drivers that are required by Generic Connectivity agents are not supplied by Oracle. Users should obtain drivers from other vendors.

The functionality of Generic Connectivity is more limited than that of Oracle Transparent Gateways.

Oracle also offers asynchronous information integration solutions that are mentioned here, but that are not discussed in this book. Briefly, these solutions include:

- Oracle Streams

Oracle Streams enables the propagation of data, transactions and events in a single data stream or queue, either within a database, or between multiple databases. Not only can Oracle Streams capture, propagate, and apply changes to data, it can also handle data structure changes (DDL) and user-defined events. Changes can be captured and applied as is, or transformed at any point in the capture, propagation, and apply processing.

- Messaging Gateway

The messaging gateway enables communication between Oracle and other non-Oracle message queuing.

- Open System Interfaces

Oracle offers a number of open interfaces, such as OCI, JDBC, and ODBC, that enable customers to use third party applications or to write their own client applications to access the Oracle database.

Benefits of Oracle's Solution for Synchronous Information Integration

Much of the processing power of Generic Connectivity and Transparent Gateways is integrated into the database. This provides an efficient solution for information integration that enables full exploitation of the power and features of the Oracle database. This includes such features as powerful SQL parsing and distributed optimization capabilities.

The following sections explore the benefits of Oracle's approach to resolving the challenges of a heterogeneous environment:

- Remote Data Can Be Accessed Transparently
- There is No Unnecessary Data Duplication
- SQL Statements Can Query Several Different Databases
- Oracle's Application Development and End User Tools Can Be Used
- Users Can Talk to a Remote Database in its Own Language

Remote Data Can Be Accessed Transparently

Both Generic Connectivity and Oracle Transparent Gateways provide the ability to transparently access data in non-Oracle databases from an Oracle environment. You can create synonyms for the objects in a non-Oracle database and refer to them without having to specify a physical location. This transparency eliminates the need for application developers to customize their applications to access data from different non-Oracle systems, thus decreasing development efforts and increasing the mobility of the application.

Instead of requiring applications to interoperate with non-Oracle systems using their native interfaces (which can result in intensive application-side processing), applications can be built upon a consistent Oracle interface for both Oracle and non-Oracle systems.

There is No Unnecessary Data Duplication

Generic Connectivity and Oracle Transparent Gateways provide applications direct access to data in non-oracle databases. This eliminates the need to upload and download large amounts of data to different locations, thus reducing data duplication and saving disk storage space. Also, by eliminating this need to upload and download large amounts of data there is a further benefit of a reduced risk for unsynchronized or inconsistent data.

SQL Statements Can Query Several Different Databases

The Oracle database server accepts SQL statements that query data stored in several different databases. The Oracle database server with the Heterogeneous Services component 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.

Oracle's Application Development and End User Tools Can Be Used

Generic Connectivity and Oracle Transparent Gateways extend the range of Oracle's database and application development tools. Oracle has tools that increase application development and user productivity by reducing prototype, development, and maintenance time.

You are not required to develop new tools or learn how to use other tools to access data stored in non-Oracle databases. Instead, you can access Oracle and non-Oracle data with a single set of Oracle tools. These tools can run on remote machines connected through Oracle Net to the Oracle database server.

Users Can Talk to a Remote Database in its Own Language

Oracle enables you 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, Oracle has a pass-through feature that enables you to bypass Oracle's query processor and to talk to the remote database in its own language

The Role of the Heterogeneous Services Component

Oracle's synchronous solutions for operating in a heterogeneous environment are Oracle Transparent Gateways and Generic Connectivity. The common component of the Oracle database server for supporting these solutions is Heterogeneous Services. This chapter describes the architecture and functionality of the Heterogeneous Services component and its interaction with Oracle Transparent Gateways and Generic Connectivity.

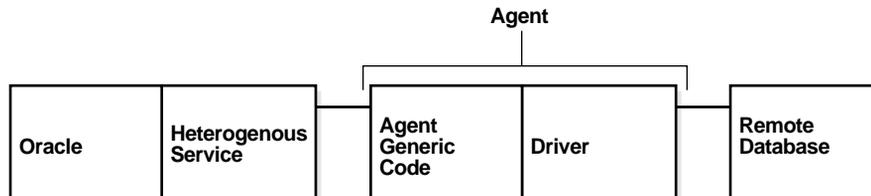
This chapter contains these topics:

- [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 component in the Oracle database server 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 component 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

A Heterogeneous Service agent is the process through which an Oracle server connects to a non-Oracle system. This agent process that accesses a non-Oracle system is called a gateway. Access to all gateways goes through the Heterogeneous Services component in the Oracle server and all gateways contain the same agent-generic code. Each gateway has a different driver linked in that maps the Heterogeneous Services application programming interface (API) to the client API of the 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: [Chapter 5, "Multithreaded Agents"](#)

Types of Heterogeneous Services Agents

There are two types of Heterogeneous Services agents:

- [Oracle Transparent Gateways](#)
- [Generic Connectivity Agents](#)

Oracle Transparent Gateways

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 being required 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 Agents

In addition to Oracle Transparent Gateways for 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 access a specific non-Oracle system using Generic Connectivity agents, you must configure an ODBC or OLE DB driver to the 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 agents have some limitations. Connections to a Generic Connectivity agent can be made only for the local Oracle database server. Functionality of these agents, especially when compared to the Oracle Transparent Gateways, is limited.

See Also: [Chapter 7, "Generic Connectivity"](#)

Heterogeneous Services Components

This section discusses the components of Heterogeneous Services in the Oracle database server. These components are:

- [Transaction Service](#)
- [SQL Service](#)

Transaction Service

The transaction service component of the Heterogeneous Services component 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 is also closed.

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 Structured Query Language (SQL) service handles the processing of all SQL-related operations. The work done by the SQL service includes:

- 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
- Translating SQL statements from Oracle's SQL dialect to the SQL dialect of the non-Oracle system
- Translating queries that reference Oracle data dictionary tables to queries that extract the necessary information from the non-Oracle system data dictionary
- Converting data from non-Oracle system datatypes to Oracle datatypes and back
- 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](#)
- [Heterogeneous Services 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.

See Also: [Appendix C, "Data Dictionary Translation Support"](#)

Heterogeneous Services Initialization Parameters

Heterogeneous Services 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 component.
- 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 Heterogeneous Services 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 datatypes 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.

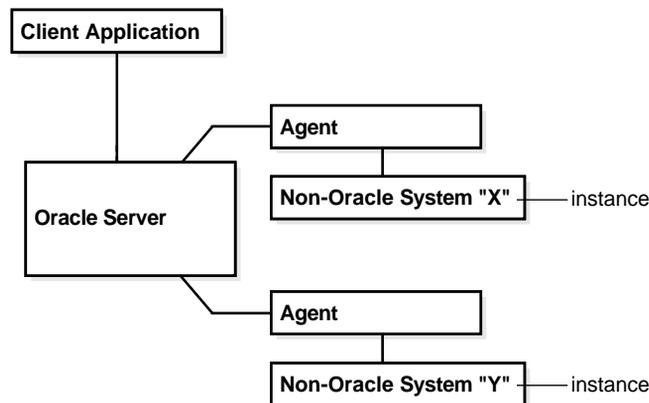
The Heterogeneous Services Data Dictionary

As mentioned in the previous section, configuration information is uploaded from an agent to the Heterogeneous Services component immediately after the connection to the agent has been established. Because 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](#), which shows two non-Oracle systems being accessed.

Figure 2-2 Accessing Multiple Non-Oracle Instances



Both agents upload configuration information, which is stored as part of the Heterogeneous Services data dictionary information on the Oracle database server.

Although it is possible to store data dictionary information at one level of granularity by having completely separate definitions in the Heterogeneous Services 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 Heterogeneous Services data dictionary 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. Each class definition (one level of granularity) is shared by all the particular instances (a second level of granularity) under that class. Further, instance information takes precedence over class information, and class information takes precedence over server-supplied defaults.

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.

Note that instance level capability and data dictionary information are session specific and hence are not stored in the Heterogeneous Services data dictionary of the Oracle database server. However, instance level initialization parameters can be stored in the database.

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:

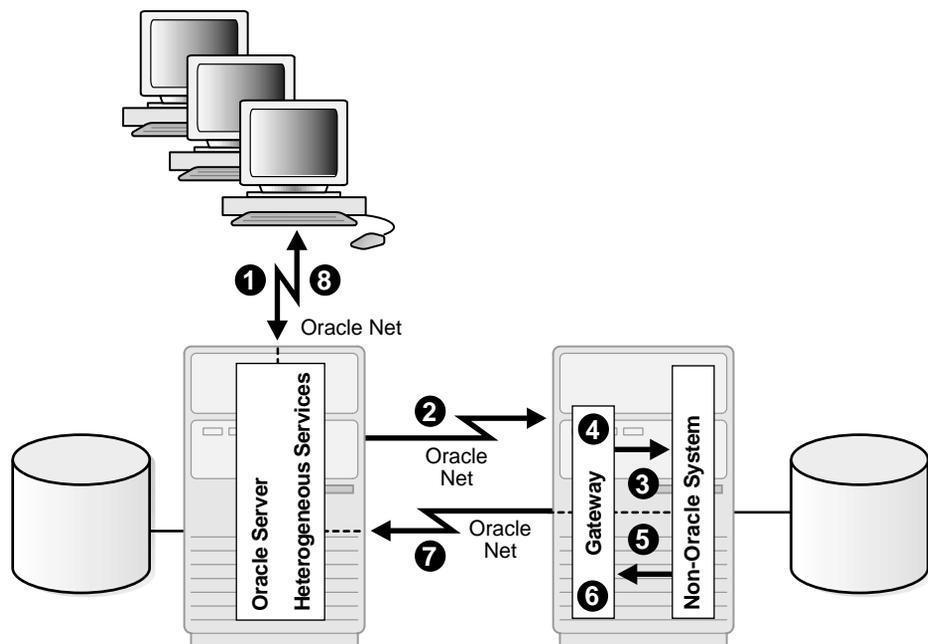
- General views
- Views used for the transaction service
- Views used for the SQL service

See Also: [Appendix C, "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 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.

Features of Oracle Transparent Gateways and Generic Connectivity

This chapter describes the major features provided by Oracle Transparent Gateways and Generic Connectivity. Descriptions of these features are contained in the following topics:

- [SQL and PL/SQL Support](#)
- [Heterogeneous Replication](#)
- [Pass-Through SQL](#)
- [Result Set Support](#)
- [Data Dictionary Translations](#)
- [Datetime Datatypes](#)
- [Two-Phase Commit Protocol](#)
- [Piecewise Long](#)
- [SQL*Plus DESCRIBE Command](#)
- [Constraints on SQL in a Distributed Environment](#)
- [Using Oracle's Optimizer with Heterogeneous Services](#)

Note: These features may not be 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 particular feature is supported for your gateway.

SQL and PL/SQL Support

SQL statements are translated and datatypes 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:

- There is no support for `CONNECT BY` clauses in SQL statements.
- ROWID support is limited; consult individual gateway documentation for more details. The Oracle Universal ROWID datatype is not supported in any gateway that uses Heterogeneous Services.
- LOBs, ADTs, and REFs are not supported.
- Remote packages are not supported.
- Remote stored procedures can have `out` arguments of type `REF CURSOR` but not `in` or `in-out` objects.
- Oracle Heterogeneous Services agents do not support shared database links.

Note: In addition to these generic limitations, 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 using materialized views.

Note: Starting with Oracle9i, Release 2, there is another means of sharing information between databases. This functionality is called Streams and includes the replication of information between Oracle and non-Oracle databases.

For information about using Streams, see *Oracle Streams Concepts and Administration*.

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 datatypes in table columns mapped to a fixed view must be compatible with (that is, have a mapping to or from) Oracle datatypes. This is usually true for datatypes 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.

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.

Modify these examples for your environment. Do not try to execute them as they are written.

Example 1: Create materialized views for heterogeneous replication

This example creates three materialized views that are then used in succeeding examples.

1. Create a primary key materialized view of table `customer@remote_db`.

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

2. Create a subquery materialized view of tables `orders@remote_db` and `customer@remote_db`.

```
CREATE MATERIALIZED VIEW sq_mv REFRESH COMPLETE AS
  SELECT * FROM orders@remote_db o WHERE EXISTS
    (SELECT c."c_id" FROM 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 customer@remote_db c,
          orders@remote_db o,
          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 heterogeneous replication

```
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;
/
```

See Also: *Oracle Database Advanced Replication* for a full description of materialized views and replication facilities

Pass-Through SQL

The pass-through SQL feature enables you to send a statement directly to a non-Oracle system without being interpreted by the Oracle 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:

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

Using the DBMS_HS_PASSTHROUGH Package

You can execute pass-through SQL statements directly at the non-Oracle system using the PL/SQL package `DBMS_HS_PASSTHROUGH`. Any statement executed with this 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: *PL/SQL Packages and Types 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 following table shows the functions and procedures provided by the DBMS_HS_PASSTHROUGH package that enable 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

Procedure/Function	Description
EXECUTE_NON_QUERY	Executes non-query
EXECUTE_IMMEDIATE	Executes non-query without bind variables
FETCH_ROW	Fetches rows from query
GET_VALUE	Retrieves column value from SELECT statement or retrieves OUT bind parameters

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 dept1 (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.

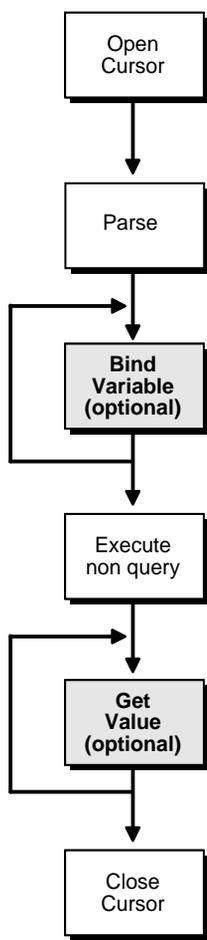
Using Bind Variables: Overview Bind variables allow you to use the same SQL statement multiple times with different values, reducing the number of times a SQL statement needs to be parsed. For example, when you need to insert four rows in a particular table, you can parse the SQL statement once and bind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

```
...
UPDATE emp
SET sal=sal*1.1
WHERE ename=:ename;
...
```

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

```
...
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 preceding 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@salesdb(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.

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

Result Set Support

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.

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 datatype called a `REF CURSOR`. Like every other Oracle datatype, a stored procedure can take this datatype 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.

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. 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 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).

See Also: Your gateway manual for more information about how result sets are supported through the gateway

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 ALL_CATALOG Oracle data dictionary view 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 in 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 wants information from the table ALL_CATALOG at Sybase, then all it has to do is to send a query referencing ALL_CATALOG@database_link to a gateway. Heterogeneous Services translates this query to the appropriate one on systables and sends 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"@remote_db SU, "dbo"."sysobjects"@remote_db 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 preceding example. 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.

Datetime Datatypes

Oracle has five datetime datatypes:

- `TIMESTAMP`
- `TIMESTAMP WITH TIME ZONE`
- `TIMESTAMP WITH LOCAL TIME ZONE`
- `INTERVAL YEAR TO MONTH`
- `INTERVAL DAY TO SECOND`

Heterogeneous Services generic code supports Oracle datetime datatypes in SQL and stored procedures. Oracle does not support these datatypes 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 by setting `HS_TIME_ZONE`.

See Also: *Oracle Database SQL Reference* for information on datetime datatypes

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: *Oracle Database Administrator's Guide* for more information about the two-phase commit protocol

Piecewise Long

Earlier versions of gateways had limited support for the `LONG` datatype. `LONG` is an Oracle datatype 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 datatype'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 example, using Heterogeneous links, you still cannot describe packages, sequences, synonyms, or types.

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 *Oracle 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 attempts 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 where the statement was issued, 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 non-Oracle databases that are accessed through Oracle Transparent Gateways or Generic Connectivity agents.

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 directly create database objects in a non-Oracle database using SQL.

However, there is an indirect way using pass-through SQL.

Consider the following example:

```
DECLARE
    num_rows INTEGER;
BEGIN
    num_rows := DBMS_HS_PASSTHROUGH.EXECUTE_IMMEDIATE@remote_db
        (
            'create table x1 (c1 char, c2 int)'
        );
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 `emp` table to a remote Oracle `emp1` table, use the following statement:

```
INSERT INTO emp@remote_db SELECT * FROM emp;
```

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 `emp`. Such a remote link received by the remote database is called a callback link.

Note: Even though callback links are supported in generic Heterogeneous Services, they may not be implemented in all Heterogeneous Services agents. Please refer to your transparent gateway documentation to determine if callback links work with the transparent gateway that you are using.

If callback links are not supported by a particular gateway, then the previous INSERT statements returns the following error:

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

The workaround is to write a PL/SQL block:

```
DECLARE
CURSOR remote_insert IS SELECT * FROM emp;
BEGIN
  FOR rec IN remote_insert LOOP
    INSERT INTO emp@remote_db (empno, ename, deptno) VALUES (
      rec.empno,
      rec.ename,
      rec.deptno
    );
  END loop;
END;
/
```

Another special case involves 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 emp@remote_db WHERE hiredate > 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 emp@remote_db WHERE hiredate > :1;
```

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

Currently, these 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, then 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 `emp1` is issued as `SELECT * FROM emp1` (with the `*` expanded to the full column list). Access for indexes is converted back to `WHERE` clause predicates. 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: [Example: 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 emp@remote_db WHERE hiredate < sysdate;
```

The statement returns the following output:

```
COUNT(*)
-----
          14
1 row selected.
```

The remote table scan is:

```
SELECT hiredate FROM emp;
```

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. Thus `sysdate` must be executed locally.

Note: Because the remote table scan operation is only partially related to the original query, the number of rows retrieved can be significantly larger than expected 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 because table expressions are not supported in the heterogeneous access module.

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

For example, consider the following type of statement:

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

This type of 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 type of 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 emp@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, then 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, then 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 workaround for this restriction is to use unique bind variables and bind by number.

Updates, Inserts, and Deletes

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
  CURSOR c1 IS SELECT empno FROM emp e, dept d
                WHERE e.deptno = d.deptno
                AND d.dname = 'SALES';
BEGIN
  FOR REC IN c1 LOOP
    UPDATE emp@remote_db SET comm = .1 * sal
      WHERE empno = rec.empno;
  END loop;
END;
/

```

Using Oracle's Optimizer with Heterogeneous Services

Oracle's optimizer can be used with Heterogeneous Services. 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.

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

Example: Using Index and Table Statistics

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.

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@remote_db;
```

This fetches all of 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 for each value of C1 in the local `t1`:

```
...  
SELECT C1 FROM remote_t1@remote_db WHERE C1 = ?;  
...
```

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

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

Enter the following:

```
EXPLAIN PLAN FOR SELECT a.* FROM remote_t1@remote_db a, T1 b  
WHERE a.C1 = b.C1;
```

Execute 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, enter the following statement:

```
SELECT ID, OTHER FROM PLAN_TABLE WHERE OPERATION = 'REMOTE';
```

Example: 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

You should see output similar to the following:

PLAN_TABLE_OUTPUT

```
-----
| Id   | Operation                                | Name | Rows  | Bytes | Cost |
| Inst | IN-OUT |
-----
|  0   | SELECT STATEMENT                          |      |      |      |      |
| * 1  | HASH JOIN                                  |      | 2000  | 197K  | 205  |
205 |
|  2   | TABLE ACCESS FULL                        | DEPT | 21    | 462   | 2    |
| * 3  | HASH JOIN                                  |      | 2000  | 154K  |      |
201 |
|  4   | REMOTE                                     |      | 2000  | 66000 |      |
|    52 |
|  5   | REMOTE                                     |      | 2000  | 92000 |      |
|    52 |
-----
```

```
PLAN_TABLE_OUTPUT
```

```
-----  
Query Block Name / Hint Alias (identified by operation id):  
-----
```

```
1 - sel$1 / D  
2 - sel$1 / D  
3 - sel$1 / F  
4 - sel$1 / F  
5 - sel$1 / E
```

```
Predicate Information (identified by operation id):
```

```
PLAN_TABLE_OUTPUT
```

```
-----  
-----  
1 - access("E"."DEPTNO"="D"."DEPTNO")  
3 - access("E"."MGR"="F"."EMPNO" AND "E"."EMPNO"="F"."EMPNO")
```

Issue the following statement:

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

You should see output similar to the following:

```
OTHER
```

```
-----  
SELECT "EMPNO", "ENAME", "DEPTNO" FROM "EMP"  
SELECT "EMPNO", "ENAME", "MGR", "DEPTNO" FROM "EMP"  
SELECT "EMPNO", "ENAME", "DEPTNO" FROM "EMP"  
SELECT "EMPNO", "ENAME", "MGR", "DEPTNO" FROM "EMP"
```

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 Heterogeneous Services Agents

This chapter explains how to use Heterogeneous Services (HS) agents. It contains the following sections:

- [Setting Up Access to Non-Oracle Systems](#)
- [Setting Initialization Parameters](#)
- [Optimizing Data Transfers Using Bulk Fetch](#)
- [Registering Agents](#)
- [Oracle Database Server SQL Construct Processing](#)
- [Executing User-Defined Functions on a Non-Oracle Database](#)
- [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: Configure Oracle Net Services to Access Heterogeneous Services Agents](#)

[Step 2: Create the Database Link to the Non-Oracle System](#)

[Step 3: Test the Connection](#)

Step 1: Configure Oracle Net Services to Access Heterogeneous Services Agents

To initiate a connection to the non-Oracle system, the Oracle server starts an agent process through the Oracle Net listener. For the Oracle 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 Oracle 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 ensure the connection uses Oracle Heterogeneous Services. 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.

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

```
Sybase_sales= (DESCRIPTION=
                (ADDRESS=(PROTOCOL=tcp)
                    (HOST=dlsun206)
                    (PORT=1521)
                )
                (CONNECT_DATA = (SERVICE_NAME=SalesDB)
                )
                (HS = OK)
            )
```

2. Set up the listener on the gateway to listen for incoming request from the Oracle server and spawn Heterogeneous Services agents. Then, start the listener on the gateway machine.

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/9.0.1)
                (PROGRAM=tg4mb80)
                (ENVS=LD_LIBRARY_PATH=non_oracle_system_lib_directory)
            )
  )

```

The value associated with the `PROGRAM` keyword defines the name of the agent executable. The full path of the directory which contains the DLL that is loaded by the Heterogeneous Services agent is specified by `LD_LIBRARY_PATH`. Typically, you use `SID_NAME` to define the initialization parameter file for the agent.

See Also:

- *Oracle Net Services Administrator's Guide* for more information about configuring the Oracle Net listener and a net service name for heterogeneous connectivity
- ["Administering Multithreaded Agents"](#) on page 5-6 for information about starting multithreaded agents

Step 2: 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` statement is the Oracle Net service name.

Use quotes with the username and password to avoid differences in case-sensitive behavior between Oracle and non-Oracle databases. For example, enter a statement like the following:

```
CREATE DATABASE LINK sales
CONNECT TO "sales1"
IDENTIFIED BY "Sales1"
USING 'Sybase_sales';
```

Step 3: 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:

```
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 Oracle data dictionary automatically the first time they are accessed. Some agent vendors may provide scripts, however, that you must run on the Oracle server.

See Also: [Heterogeneous Services Data Dictionary Views](#) on page 4-20 and [Appendix C, "Data Dictionary Translation Support"](#)

Setting Initialization Parameters

As mentioned in "[Configuring Heterogeneous Services](#)" on page 2-5, you 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 initialization 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 database server initialization parameter file (`init.ora` file). There is a Heterogeneous Services initialization file for each gateway instance.

Name and Location of Heterogeneous Services Initialization Parameter File

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 Heterogeneous Services initialization file is located in the directory `$ORACLE_HOME/hs/admin`. In the case of Transparent Gateways it is located in the directory `$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`.

Syntax for Initialization Parameter Settings

The initialization 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 = 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.

In the initialization parameter syntax, all keywords (`SET`, `PRIVATE` and `IFILE`) are case insensitive. Initialization parameter names and values are case sensitive. Most initialization parameter names are uppercase. String values for Heterogeneous Services parameters must be lowercase. Exceptions to this rule are explicitly noted.

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

IFILE = path name for file to be included

Gateway Initialization Parameters

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 following list of generic initialization parameters are the only ones discussed in this document:

HS_CALL_NAME
HS_COMMIT_POINT_STRENGTH
HS_DB_DOMAIN
HS_DB_INTERNAL_NAME
HS_DB_NAME
HS_DESCRIBE_CACHE_HWM
HS_FDS_CONNECT_INFO
HS_FDS_DEFAULT_SCHEMA_NAME
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_NUMERIC_CHARACTERS
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 prevent the parameter from being uploaded to the server and could cause errors in SQL processing. None of these parameters are required to be set in the environment, so the `SET` keyword need not be used.

See Also:

- [Appendix A, "Heterogeneous Services Initialization Parameters"](#) for descriptions of the generic Heterogeneous Services initialization parameters
- Individual gateway documentation for the list of initialization parameters specific to a gateway

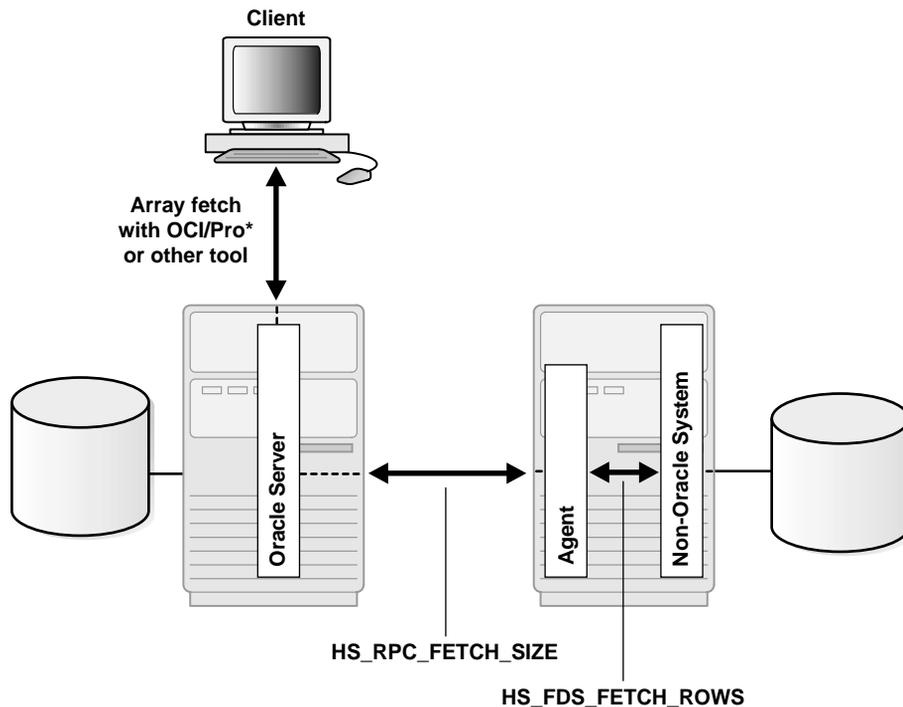
Optimizing Data Transfers Using Bulk Fetch

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

1. From the non-Oracle system to the agent process
2. From the agent process to the Oracle database server
3. From the Oracle database server to the application

Oracle optimizes 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 Oracle 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 or each 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 for each fetch between the agent and the Oracle 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 Oracle 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 Oracle9i and later, 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.
Capability definitions	Identify details like SQL language features supported by the non-Oracle data source.
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 interoperability between any Oracle database server and any Heterogeneous Services agent, if 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 Oracle 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 for agent self-registration 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 4-20 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 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: *Oracle 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:

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

The non-Oracle 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"@remote_db;
```

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"@remote_db 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 side, so an error is issued.

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

Executing User-Defined Functions on a Non-Oracle Database

You can execute user-defined functions in a remote non-Oracle database. For example:

```
SELECT getdeptforemp@Remote_DB(7782) FROM dual;
```

In this example, the user issues a `SELECT` statement that executes a user-defined function in the remote database that returns department information for employee 7782.

When the remote function resides in an Oracle database, then the Oracle database automatically ensures that the remote function does not update any database state (such as updating rows in a database or updating the PL/SQL package state). The gateway cannot verify this when the remote function resides in a non-Oracle database. Therefore, the user is responsible for ensuring that the user-defined functions do not update the state in any database. Ensuring no updates to the database is required to guarantee read consistency.

As a security measure, you must specify the functions that you want to execute remotely and their owners in the `HS_CALL_NAME` parameter in the gateway-specific initialization parameter file. For example:

```
HS_CALL_NAME = "owner1.A1, owner2.A2, "
```

`owner1` and `owner2` are the remote function owner names. `A1` and `A2` are the remote function names. You do not need to specify the remote function owner in the SQL statement. By default, the remote function needs to reside in the schema that the Transparent Gateway connects to. If this is not the case, then you must specify the owner of the remote function in the SQL statement.

Some other examples of executing user-defined remote functions are: as follows:

- A remote function in a subquery

The function uses the `employee_id` column data to retrieve the `department_id` from the `EMPLOYEES` table in the remote database. The outer query then determines all department numbers in the remote database that match the returned list.

```
SELECT * FROM departments@remotedb
WHERE department_id IN
  (SELECT
   getdeptforemp@remotedb(employee_id)
   FROM employees@remotedb);
```

- Applying a local function to the result of a user-defined remote function

This query returns the maximum salary of all employees on the remote database.

```
SELECT max(getsalforemp@remotedb(employee_id))
FROM employees@remotedb;
```

- A DML statement

The statement uses the output from a user-defined query in the remote database to update the salary column with new salary information.

```
UPDATE employee_history
SET salary = emp_changed_salary@remote_db;
```

In these examples, the Oracle database passes the function name and owner to the Transparent Gateway. The user-defined function is executed on the remote database.

See Also: ["HS_CALL_NAME"](#) on page A-3

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 table.

The following statement creates a system wide synonym for the `emp` table in the schema of user `ORACLE` in the Sybase database:

```
CREATE PUBLIC SYNONYM emp FOR "ORACLE"."EMP"@SYBS;
```

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

Example of a Distributed Query

Note: Modify these examples for your environment. Do not try to execute them as they are written.

The following statement joins data between the Oracle database server, an IBM DB2 database, and a Sybase database:

```
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:

```
CREATE SYNONYM ORDERS FOR ORDERS@DB2;
CREATE SYNONYM PROJECTS FOR "PROJECTS"@SYBS;
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 statement:

```
SELECT * FROM DETAILS;
```

The statement 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

In Oracle9i, release 2 and later, Heterogeneous Services supports callback links. This enables SQL statements like the following to be executed:

```
INSERT INTO table_name@dblink SELECT column_list FROM table_name;
```

Even though Heterogeneous Services supports the callback functionality, not all gateways have implemented it. If the gateway that you are using has not implemented this functionality, then the preceding INSERT statement returns the following error message:

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

See Also: Your gateway documentation for information about support for callback links

For gateways that do not support callback links, you can use the SQL*Plus COPY command. 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, and commits the transaction:

```
COPY FROM SCOTT/TIGER@inst1 -
INSERT EMP@remote_db -
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. 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

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

The `CREATE TABLE` statement lets you copy data from a non-Oracle database to the Oracle database server. To create a table on the local database and insert rows from the non-Oracle 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:

```
CREATE TABLE table1 AS SELECT * FROM "EMP"@remote_db;
```

Alternatively, you can use the SQL*Plus `COPY` command to copy data from the non-Oracle 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
<code>HS_BASE_CAPS</code>	SQL service	All capabilities supported by Heterogeneous Services
<code>HS_BASE_DD</code>	SQL service	All data dictionary translation table names supported by Heterogeneous Services
<code>HS_CLASS_CAPS</code>	Transaction service, SQL service	Capabilities for each class
<code>HS_CLASS_DD</code>	SQL service	Data dictionary translations for each class

Table 4–1 Data Dictionary Views for Heterogeneous Services

View	Type	Identifies
HS_CLASS_INIT	General	Initialization parameters for each class
HS_FDS_CLASS	General	Classes accessible from the Oracle server
HS_FDS_INST	General	Instances accessible from the Oracle server

Like all Oracle data dictionary tables, these views are read-only. Do not change the content of any of the underlying tables.

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 Oracle 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:

Type	Capability
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 cannot 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 cannot 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 cannot 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 LIKE 'SYBASE%';
```

```
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 Oracle 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 Oracle 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@remote_db
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, connect as user `SYS` and issue the following query on the `HS_CLASS_DD` view:

```
SELECT DD_TABLE_NAME, TRANSLATION_TYPE
FROM   HS_CLASS_DD
WHERE  FDS_CLASS_NAME LIKE 'SYBASE%';
```

```
DD_TABLE_NAME          T
-----
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 C, "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 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
V\$HS_AGENT	Identifies the set of Heterogeneous Services agents currently running on a given host, using one row for each 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 *Oracle Database Reference*):

Table 4–2 V\$HS_AGENT

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

Table 4–3 V\$HS_SESSION

Column	Description
DB_LINK	Server database link name used to access the agent NULL means that no database link is used (for example, 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 *Oracle 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 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

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:

- [Why Use Multithreaded Agents?](#)
- [Multithreaded Agent Architecture](#)
- [Administering Multithreaded Agents](#)

Note: Heterogeneous Services supports multithreaded agents, however this functionality may not be available in all Heterogeneous Services based gateways. In addition to the generic support for multithreaded agents that Heterogeneous Services provides, multithreaded agents support must be added to the driver. Please refer to your transparent gateway documentation to determine if the gateway you are using is certified to work as a multithreaded agent.

See Also: *Oracle Database Application Developer's Guide - Fundamentals* for information on how the `extproc` agent can be configured to run in MTA mode

Why Use Multithreaded Agents?

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

By default, a Heterogeneous Services agent is started up for each user session. 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 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 for each one of them, there are several thousand agent processes running concurrently. 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 a disproportionate amount of system resources without any discernible benefit.

In the case of connections to the Oracle database server, this problem is addressed by starting the server in shared server mode. Shared server mode allows database connections to be shared by a small number of server processes.

The Advantage of Multithreading

The Oracle shared server architecture assumes that even when there are several thousand user sessions currently open, only a small percentage of these connections will be active at any given time. In shared server mode, there is a pool of shared

server processes. User sessions connect to dispatcher processes that place the tasks requested by the user sessions on a queue. The tasks are picked up by the first available shared server processes. The number of shared server processes is usually considerably less than the number of user sessions.

Multithreaded Heterogeneous Services agents provide similar functionality for connections to non-Oracle systems. The multithreaded agent architecture uses a pool of shared agent threads. The tasks requested by the user sessions are put on a queue and are picked up by the first available multithreaded agent thread. Because only a small percentage of user connections are actually active at a given moment, using a multithreaded architecture allows for more efficient use of system resources.

Multithreaded Agent Architecture

Multithreaded agents must be prestarted on a one for each system identifier (SID) basis. This is done using the agent control utility `agctl`. This utility is also used to configure the agent and to shut down the agent.

Each TNS listener that is running on a system listens for incoming connection requests for a set of SIDs. If the SID in an incoming Oracle Net connect string is one of the SIDs that the listener is listening for, then that listener will process the connection. Further, if a multithreaded agent has been started for the SID, then the listener will pass the request to that agent.

In the architecture for multithreaded agents, each incoming connection request is processed by means of the three different kinds of threads:

- A single **monitor** thread

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

- Several **dispatcher** threads

The dispatcher threads are responsible for the following:

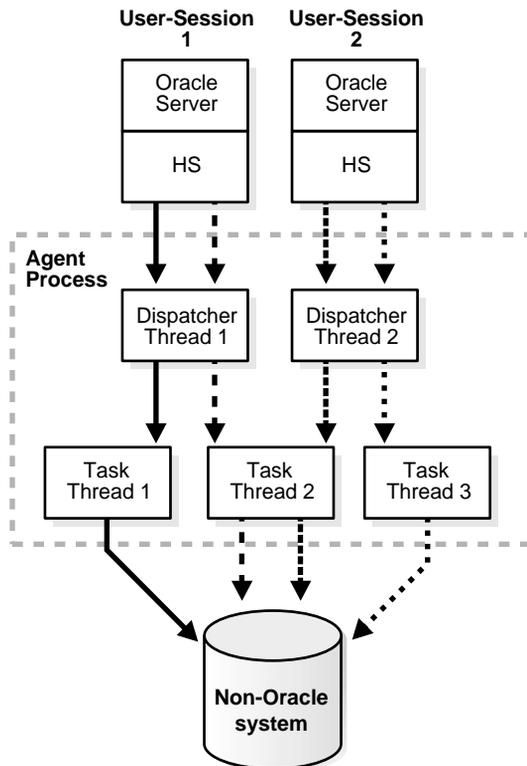
- Handling communication with the Oracle server
- Passing task requests on to the task threads

- Several **task** threads

The task threads handle requests from the Oracle processes.

The multithreaded agent architecture is illustrated in [Figure 5-1](#) where each request issued by a user session is represented in by a separate type of arrow. There is no representation of the monitor thread in this illustration, because that thread is created once when the multithreaded agent is started and it creates and monitors the other threads. Typically there are many more task threads than dispatcher threads.

Figure 5-1 Multithreaded Agent Architecture



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.

Each type of thread is discussed in more detail in the following sections:

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

See Also: [Administering Multithreaded Agents](#) on page 5-6 for more information on how to start and stop multithreaded agents using the agent control utility

The Monitor Thread

When a multithreaded agent is started for a SID by the agent control utility, the monitor thread is created. The monitor thread performs the following functions:

- It creates the dispatcher and task threads.
- It 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.

- It monitors the other threads and sends load information about the dispatcher threads to all the listener processes handling connections to this agent.

This enables the listeners to hand over incoming connections to the least loaded dispatcher.

- It continues to monitor each of the threads it has created.

Dispatcher Threads

Dispatcher threads perform the following functions:

- They accept incoming connections and task requests from Oracle servers.

- They place incoming requests on a queue for a task thread to pick up.
- They 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 perform the following functions:

- They pick up requests from a queue.
- They perform the necessary operations.
- They place the results on a queue for a dispatcher to pick up.

Administering Multithreaded Agents

As discussed earlier, multithreaded Heterogeneous Services agents must be prestarted on a one for each system identifier (SID) basis before any attempt is made to connect to the non-Oracle system. Any agent not spawned in this fashion will not function in multithreaded mode, and must be set up as described in "[Setting Up Access to Non-Oracle Systems](#)" on page 4-2.

A multithreaded agent is started, stopped, and configured by an agent control utility called `agtctl`, which works much like `lsnrctl`. However, unlike `lsnrctl`, which reads a configuration file (`listener.ora`), `agtctl` takes configuration information from the command line and writes it to a control file.

The following topics are discussed in this section:

- [Agent Control Utility \(agtctl\) Commands](#)
- [Using Single-Line Command Mode](#)
- [Using Shell Mode Commands](#)
- [Configuration Parameters for Multithreaded Agent Control](#)

Agent Control Utility (agtctl) Commands

You start and stop `agtctl`, and create and maintain its control file, using the commands shown in [Table 5-1](#).

Table 5-1 Agent Control Utility Commands

Command	Description
<code>startup</code>	Starts a multithreaded agent
<code>shutdown</code>	Stops a multithreaded agent
<code>set</code>	Sets a configuration parameter for a multithreaded agent
<code>unset</code>	Causes a parameter to revert to its default value
<code>show</code>	Displays the value of a configuration parameter
<code>delete</code>	Deletes the entry for a particular SID from the control file
<code>exit</code>	Exits shell mode
<code>help</code>	Lists available commands

These commands can be issued in one of two ways:

- You can issue commands from the UNIX (or DOS) shell.
This mode is called single-line command mode.
- You can type `agtctl` and a "`AGTCTL>`" prompt appears. You then can type commands from within the `agtctl` shell.
This mode is called shell mode.

The syntax and parameters for `agtctl` commands vary depending upon the mode in which they are issued.

Note:

- All commands are case sensitive.
 - The agent control utility puts its control file in either the directory pointed to by the environment variable `AGTCTL_ADMIN` or in the directory pointed to by the environment variable `TNS_ADMIN`. Ensure that at least one of these environment variables is set and that it points to a directory that the agent has access to
 - If the Heterogeneous Services agent requires an environment variable to be set, or if the `ENVS` parameter was used when configuring the `listener.ora` entry for the agent working in dedicated mode, then all required environment variables must be set in the UNIX (or DOS) shell which runs the `agtctl` utility.
-

Using Single-Line Command Mode

This section describes the use of `agtctl` commands. They are presented in single-line command mode.

Setting Configuration Parameters for a Multithreaded Agent

You should set the configuration parameters for a multithreaded agent before you start the agent. They determine how the agent will be configured. If a configuration parameter is not specifically set, a default value is used. Configuration parameters and their default values are shown in [Table 5-2](#).

Use the `set` command to set multithreaded agent configuration parameters.

Syntax

```
agtctl set parameter parameter_value agent_sid
```

parameter is the parameter that you are setting.

parameter_value is the value being assigned to the parameter.

agent_sid is the SID that this agent will service. Must be specified for single-line command mode.

Example

```
agtctl set max_dispatchers 5 salesDB
```

Starting a Multithreaded Agent

Use the `startup` command to start an agent in multithreaded mode.

Syntax

```
agtctl startup agent_name agent_sid
```

agent_name is the name of the agent. For example, `extproc` is an agent name. *agent_sid* is the SID that this agent will service. Must be specified for single-line command mode.

Example

```
agtctl startup extproc salesDB
```

Shutting Down a Multithreaded Agent

Use the `shutdown` command to stop a multithreaded agent. There are three forms of shutdown.

- Normal

This form of shutdown is the default. It causes `agtctl` will talk to the agent and ask it to terminate itself gracefully. All sessions complete the operations they are currently doing and then shutdown.

- Immediate

In this form of shutdown, `agtctl` talks to the agent and tells it to terminate immediately. The agent process exits immediately regardless of the state of current sessions.

- Abort

In this form of shutdown, `agtctl` does not talk to the agent at all. It just issues a system call to kill the agent process.

Syntax

```
agtctl shutdown [immediate|abort] agent_sid
```

agent_sid is the SID which this agent will service. Must be specified for single-line command mode.

Example

```
agtctl shutdown immediate salesDB
```

Examining the Value of Configuration Parameters

To examine the value of a configuration parameter use the `show` command.

Syntax

```
agtctl show parameter agent_sid
```

parameter is the parameter that you are examining.

Example

```
agtctl show max_dispatchers salesDB
```

Resetting a Configuration Parameter to Its Default Value

You can reset a configuration parameter to its default value using the `unset` command.

Syntax

```
agtctl unset parameter agent_sid
```

parameter is the parameter that you are examining.

agent_sid is the SID which this agent will service. Must be specified for single-line command mode.

Example

```
agtctl unset max_dispatchers salesDB
```

Deleting an Entry for a Specific SID from the Control File

The `delete` command deletes the entry for the specified SID from the control file.

Syntax

```
agtctl delete agent_sid
```

agent_sid is the SID entry to delete.

Example

```
agtctl delete salesDB
```

Requesting Help

Use the `help` command to view a list of available commands for `agtctl`, or to see the syntax for a particular command.

Syntax

```
agtctl help [command]
```

command is the command whose syntax you want to view.

Example

```
agtctl help set
```

Using Shell Mode Commands

In shell mode, you start `agtctl` by typing `agtctl` whereupon you will see an "AGTCTL>" prompt. Thereafter, since you are issuing commands from within the `agtctl` shell, you do not prefix the command string with the word `agtctl`.

Next, set the name of the agent SID that you are working with by entering the following:

```
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. Unlike single-line command mode, you do not specify *agent_sid* in the command string.

You can optionally set the language for error messages, to other than English, as follows:

```
set language language
```

The commands themselves are the same as those for the single-line command mode. To exit shell mode, type `exit`.

The following are examples of shell mode commands.

Example: Setting a Configuration Parameter

This example sets a new value for the `shutdown_address` configuration parameter.

```
set shutdown_address (address=(protocol=ipc)(key=oraDBsalesDB))
```

Example: Starting a Multithreaded Agent

This example starts a multithreaded agent.

```
startup extproc
```

Configuration Parameters for Multithreaded Agent Control

The following table lists the configuration parameters for the agent control utility.

Table 5–2 Initialization Parameters for agtctl

Parameter	Description	Default Value
max_dispatchers	Maximum number of dispatchers	1
tcp_dispatchers	Number of dispatchers listening on tcp (the rest are using ipc)	0
max_task_threads	Number of task threads	2
max_sessions	Maximum number of sessions	5
listener_address	Address on which the listener is listening (needed for registration)	(ADDRESS_LIST= (ADDRESS= (PROTOCOL=IPC) (KEY=PNPKEY)) (ADDRESS= (PROTOCOL=IPC) (KEY=oracle_sid)) (ADDRESS= (PROTOCOL=TCP) (HOST=127.0.0.1) (PORT=1521)))
		Note: <i>oracle_sid</i> is the SID of the Oracle database.
shutdown_address	Address on which the agent should listen for shutdown messages from agtctl	(ADDRESS= (PROTOCOL=IPC) (KEY=oracle_sid agent_sid))
		Notes:
		<ul style="list-style-type: none"> ▪ <i>agent_sid</i> is the SID of the multithreaded agent. ▪ indicates that <i>oracle_sid</i> and <i>agent_sid</i> are concatenated into one string.

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](#)
- [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 Oracle supports.

The Oracle Transparent Gateways tell Oracle (at connect time) which functions and operators they do support. If the non-Oracle 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.

Optimizing Performance of Distributed Queries

You can improve performance of distributed queries in several ways. These are:

- 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 query optimizer.

The query optimizer uses indexes on remote tables, considers more execution plans than the rule-based optimizer, and generally gives better results. With the query 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

- Use 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 and Functions](#)
- [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 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. ODBC agents are supported on Solaris, HP-UX, AIX, and Windows NT platforms. OLE DB agents are supported only on the Windows NT platform.

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](#)
- [Datatype 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

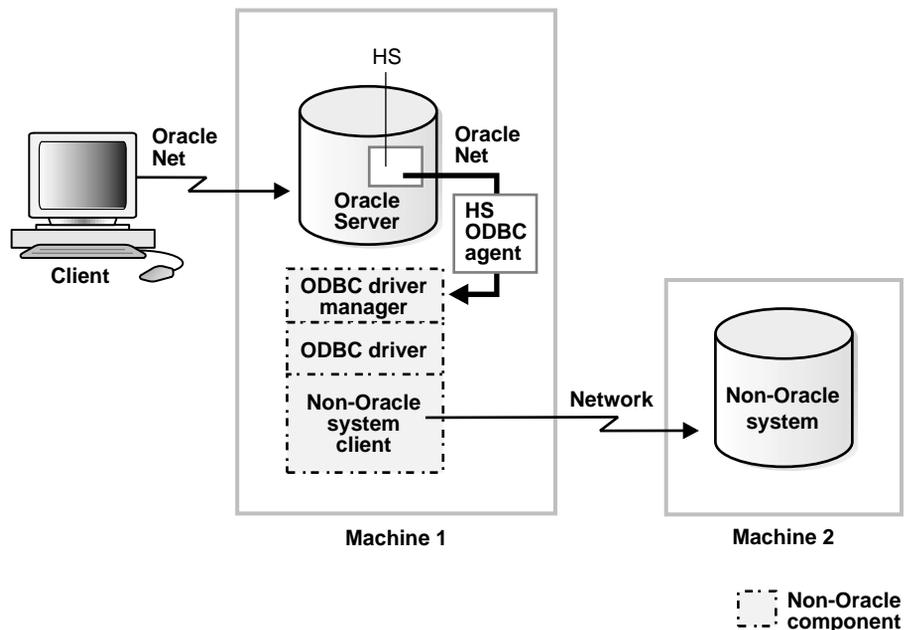
To access the non-Oracle data store using Generic Connectivity, the agents work with an ODBC or OLE DB driver. The Oracle database server provides support for

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 a 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

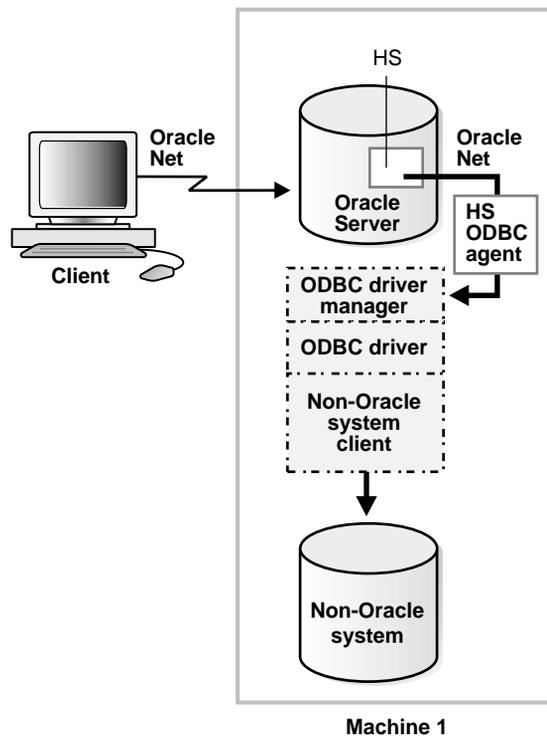
- An ODBC driver
- A non-Oracle client application

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. Please refer to your ODBC driver documentation for information about the requirements for the ODBC driver.

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.

Datatype Mapping

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

If a table contains a column whose datatype is not supported by Generic Connectivity, then the column information is not returned to the Oracle server.

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` and `CLOB` data cannot be read through passthrough queries.

- Updates or deletes that include unsupported functions within a `WHERE` clause are not allowed.
- Generic Connectivity does not support stored procedures.
- Generic Connectivity agents cannot participate in distributed transactions; they support single-site transactions only.
- Generic Connectivity does not support multithreaded agents.

Supported Oracle SQL Statements and Functions

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
```

Generic Connectivity assumes that the following minimum set of SQL functions is supported by the ODBC driver or OLE DB provider that is being used:

- `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, `inithsodbc.ora` for ODBC agents and `inithsoledb.ora` for OLE DB agents. 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: ["Setting Initialization Parameters"](#) on page 4-5 for more information on parameters

Set the parameter values as follows:

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

[SET] and [PRIVATE] are optional keywords. If you do not specify either SET or PRIVATE, then 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 ["Setting Initialization Parameters"](#) on page 4-5 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

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.

Specify the value as follows:

```
HS_FDS_CONNECT_INFO=dsn
```

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, version 3.5:

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. Refer to your ODBC driver documentation and follow the prompts to complete configuration of the DSN.
8. After creating the system DSN, click **OK** to 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_path_of_odbc_driver
```

If your ODBC driver requires an ODBC driver manager, then `HS_FDS_SHAREABLE_NAME` should point to the location of the ODBC driver manager and not to the 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 "[Setting Initialization Parameters](#)" on page 4-5 for detailed information on all the initialization parameters.

Setting Parameters on UNIX: Example The following is an example of an `odbc.ini` file that uses DataDirect Technologies SQLServer ODBC driver:

```
[ODBC Data Sources]
SQLServerWP=DataDirect 4.10 SQL Server Wire Protocol

[SQLServerWP]
Driver=/opt/odbc410/lib/ivmsssl8.so
Description=DataDirect 4.10 SQL Server Wire Protocol
Database=oratst
LogonID=TKHOUSER
```

```

Password=TKHOUSER
Address=sqlserver-pc,1433
QuotedId=No
AnsiNPW=No

[ODBC]
Trace=0
TraceFile=/opt/odbc410/odbctrace.out
TraceDll=/opt/odbc410/lib/odbctrac.so
InstallDir=/opt/odbc410
ConversionTableLocation=/opt/odbc410/tables
UseCursorLib=0

```

To configure the Generic Connectivity ODBC agent to use this driver, the following lines are required in `initHS_SID.ora`:

```

HS_FDS_CONNECT_INFO=SQLServerWP
HS_FDS_SHAREABLE_NAME=/opt/odbc4/lib/libodbc.so
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"

```

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 "[Setting Initialization Parameters](#)" on page 4-5 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 32-bit ODBC driver that conforms to ODBC version 2.5 is required.
- On UNIX machines, a 32-bit ODBC driver which conforms to ODBC version 2.5 is required. If the ODBC driver works with an ODBC driver manager that conforms to ODBC version 2.5, then the ODBC driver version can be 2.5 or higher.

See Also: Your ODBC driver documentation for dependencies on an ODBC driver manager

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 datatypes and should support SQL grammar level `SQL_92`. The ODBC driver should also expose the following ODBC APIs:

Table 7-1 ODBC Functions

ODBC Function	Comment
<code>SQLAllocConnect</code>	-
<code>SQLAllocEnv</code>	-
<code>SQLAllocStmt</code>	-
<code>SQLBindCol</code>	-
<code>SQLBindParameter</code>	-

Table 7-1 ODBC Functions(Cont.)

ODBC Function	Comment
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	-
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	-

Table 7-1 ODBC Functions(Cont.)

ODBC Function	Comment
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 a 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	Notes
IAccessor	CreateAccessor, ReleaseAccessor	-
IColumnsInfo	GetColumnsInfo (Command and Rowset objects)	-
ICommand	Execute	-
ICommandPrepare	Prepare	-
ICommandProperties	SetProperties	-
ICommandText	SetCommandText	-
ICommandWithParameters	GetParameterInfo	-
IDBCreateCommand	CreateCommand	-
IDBCreateSession	CreateSession	-
IDBInitialize	Initialize	-

Table 7–2 OLE DB (SQL) Interfaces

Interface	Methods	Notes
IDBSchemaRowset	GetRowset (tables, columns, indexes; optionally also procedures, procedure parameters)	-
IErrorInfo	GetDescription, GetSource	You can also use IErrorLookup with the GetErrorDescription method.
IErrorRecords	GetErrorInfo	-
ILockBytes (OLE)	Flush, ReadAt, SetSize, Stat, WriteAt	Required only if BLOB datatypes are used in the OLE DB provider.
IRowset	GetData, GetNextRows, ReleaseRows, RestartPosition	-
IStream (OLE)	Read, Seek, SetSize, Stat, Write	-
ISupportErrorInfo	InterfaceSupportsErrorInfo	-
ITransactionLocal (optional)	StartTransaction, Commit, Abort	-

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 Interfaces for Data Providers to Expose

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

Table 7–3 OLE DB (FS) Interfaces

Interface	Methods	Notes
IAccessor	CreateAccessor, ReleaseAccessor	-

Table 7–3 OLE DB (FS) Interfaces (Cont.)

Interface	Methods	Notes
IColumnsInfo	GetColumnsInfo (Command and Rowset objects)	-
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)	Flush, ReadAt, SetSize, Stat, WriteAt	Required only if BLOB datatypes are used in the OLE DB provider
IRowsetIndex	SetRange	Required only if indexes are used in the OLE DB provider
IErrorInfo	GetDescription, GetSource	You can use IErrorLookup with the GetErrorDescription method as well.
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	-

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_OWNINGINSERT = TRUE
- DBPROP_OTHERUPDATEDELETE = TRUE
- DBPROP_CANSROLLBACKWARDS = TRUE
- DBPROP_IRowsetLocate = TRUE
- DBPROP_OTHERINSERT = FALSE

Heterogeneous Services Initialization Parameters

The Heterogeneous Services initialization parameter file contains configuration settings stored as a text file.

This section contains the following topics:

- [HS_CALL_NAME](#)
- [HS_COMMIT_POINT_STRENGTH](#)
- [HS_DB_DOMAIN](#)
- [HS_DB_INTERNAL_NAME](#)
- [HS_DB_NAME](#)
- [HS_DESCRIBE_CACHE_HWM](#)
- [HS_FDS_CONNECT_INFO](#)
- [HS_FDS_DEFAULT_SCHEMA_NAME](#)
- [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_NUMERIC_CHARACTER](#)

-
- [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](#)

See Also:

- ["Setting Initialization Parameters"](#) on page 4-5 contains instructions for setting the Heterogeneous Services initialization parameters
- ["Configuring Generic Connectivity Agents"](#) on page 7-7 contains instructions for setting Heterogeneous Services initialization parameters specific to Generic Connectivity

HS_CALL_NAME

Property	Description
Default value	None
Range of values	Not applicable

Specifies the remote functions that can be referenced in SQL statements. The value is a list of remote functions and their owners, separated by semicolons, in the following format:

```
owner_name.function_name
```

For example:

```
owner1.A1;owner2.A2;owner3.A3
```

If no owner name is specified for a remote function, then the default owner name is the user name used to connect to the remote database (specified when the HS database link is created).

The entries for the owner names and the function names are case-sensitive.

See Also: ["Executing User-Defined Functions on a Non-Oracle Database"](#) on page 4-15

HS_COMMIT_POINT_STRENGTH

Property	Description
Default value	0
Range of values	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 *Oracle 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

Property	Description
Default value	WORLD
Range of values	1 to 199 characters

Specifies a unique network sub-address for a non-Oracle system. HS_DB_DOMAIN is similar to DB_DOMAIN, described in the *Oracle 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

Property	Description
Default value	01010101
Range of values	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 nonunique number can cause problems when two-phase commit recovery actions are necessary for a transaction.

HS_DB_NAME

Property	Description
Default value	HO
Range of values	1 to 8 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

Property	Description
Default value	100
Range of values	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

Property	Description
Default value	None
Range of values	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 specify a file DSN (data source name) or a system DSN as follows:

```
HS_FDS_CONNECT_INFO=FILEDSN=dsn
```

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 format:

```
HS_FDS_CONNECT_INFO="UDLFILE=data_link"
```

Note: Whenever the parameter value includes an equal sign (=), it must be enclosed in quotation marks.

HS_FDS_DEFAULT_SCHEMA_NAME

Property	Description
Default value	None
Range of values	Not applicable

Specifies a default value for the owner column that will be returned in the ddtrans, when the value of the owner is null. For example:

```
HS_FDS_DEFAULT_SCHEMA_NAME = PUBLIC
```

HS_FDS_SHAREABLE_NAME

Property	Description
Default value	None
Range of values	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. If your ODBC driver requires an ODBC driver manager, then `HS_FDS_SHAREABLE_NAME` should point to the location of the ODBC driver manager and not to the ODBC driver.

HS_FDS_TRACE_LEVEL

Property	Description
Default value	OFF
Range of values	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

Property	Description
Default value	System-specific
Range of values	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 globalization 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 Oracle 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 globalization support between the Oracle database server and the non-Oracle data source depends on how the driver is implemented.

HS_LONG_PIECE_TRANSFER_SIZE

Property	Description
Default value	64 KB
Range of values	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

Property	Description
Default value	Value determined by <code>HS_LANGUAGE</code> parameter
Range of values	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 *Oracle Database SQL 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

Property	Description
Default value	Value determined by HS_LANGUAGE parameter
Range of values	Any valid NLS_LANGUAGE 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

Property	Description
Default value	Value determined by HS_LANGUAGE parameter
Range of values	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-7

HS_NLS_NUMERIC_CHARACTER

Property	Description
Default value	Value determined by HS_LANGUAGE parameter

Property	Description
Range of values	Any valid NLS_NUMERIC_CHARACTERS value (any two valid numeric characters)

Specifies the characters to use as the group separator and the decimal character. The group separator separates integer groups (such as thousands, millions, and billions). The decimal character separates the integer portion of a number from the decimal portion.

HS_NLS_TIMESTAMP_FORMAT

Property	Description
Default value	Derived from NLS_TERRITORY
Range of values	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 *Oracle Database SQL 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

Property	Description
Default value	Derived from NLS_TERRITORY
Range of values	Any valid datetime with time zone format mask

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 *Oracle Database SQL 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

Property	Description
Default value	50
Range of values	1 to the 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

Property	Description
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

Property	Description
Default value	ON
Range of values	OFF or 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

Property	Description
Default value	4000
Range of values	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

Property	Description
Default value for '[+ -]hh:mm'	Derived from NLS_TERRITORY
Range of values for '[+ -]hh:mm'	Any valid datetime format mask

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

Property	Description
Default value	None
Range of values	Valid parameter file names

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 *Oracle Database Reference*

Datatype Mapping

Oracle maps the ANSI datatypes through ODBC and OLE DB interfaces to supported Oracle datatypes. When the results of a query are returned, Oracle9i converts the ODBC or OLE DB datatypes to Oracle datatypes.

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.

This appendix contains the following tables:

- [Mapping ANSI Datatypes to Oracle Datatypes Through an ODBC Interface](#)
- [Mapping ANSI Datatypes to Oracle Datatypes Through an OLE DB Interface](#)

Mapping ANSI Datatypes to Oracle Datatypes Through an ODBC Interface

Table B-1 maps ODBC datatypes into equivalent ANSI and Oracle datatypes. In some cases equivalence to ANSI datatypes 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 datatypes for a typical implementation of ANSI SQL.

Table B-1 Mapping ANSI Datatypes to Oracle Datatypes 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)
		Note: It is possible under some circumstance for the INTEGER ANSI datatype to map to Precision 38, but it usually maps to Precision 10.
N/A	SQL_LONGVARBINARY	LONG RAW
N/A	SQL_LONGVARCHAR	LONG
		Note: 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).
REAL	SQL_REAL	FLOAT (23)

Table B-1 Mapping ANSI Datatypes to Oracle Datatypes Through an ODBC Interface

ANSI	ODBC	Oracle
SMALLINT	SQL_SMALLINT	NUMBER (5)
TIME	SQL_TIME	DATE
TIMESTAMP	SQL_TIMESTAMP	DATE
NUMERIC (3 , 0)	SQL_TINYINT	NUMBER (3)
VARCHAR	SQL_VARCHAR	VARCHAR
BOOLEAN	SQL_BIT	RAW (1)

Mapping ANSI Datatypes to Oracle Datatypes Through an OLE DB Interface

Table B-2 Mapping ANSI Datatypes to Oracle Datatypes Through an OLE DB Interface

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

Table B–2 Mapping ANSI Datatypes to Oracle Datatypes Through an OLE DB Interface (Cont.)

ANSI	OLE DB	Oracle
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

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](#)
- [Views and Tables Supported by Generic Connectivity](#)

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 Oracle 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 Oracle data dictionary table.
4. Passes the data dictionary information from the non-Oracle system table to Oracle.

Note: The values returned when querying the Generic Connectivity data dictionary may not be the same as the ones returned by the Oracle `SQL*Plus 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_MVIEWS
ALL_SOURCE
ALL_SYNONYMS
ALL_TABLES
ALL_TAB_COLUMNS
ALL_TAB_COMMENTS
ALL_TAB_PRIVS
ALL_TAB_PRIVS_MADE
ALL_TAB_PRIVS_REC'D
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_MVIEW_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

Views and Tables Supported by Generic Connectivity

Generic Connectivity supports only the views and tables shown in [Table C-1](#).

If you use an unsupported view, then you receive an Oracle error message that no rows were 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 C-1 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

Table C-1 Generic Connectivity Data Dictionary Mapping (Cont.)

View	ODBC API	OLE DB API
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
USER_CONS_COLUMNS	SQLPrimaryKeys, SQLForeignKeys	DBSCHEMA_PRIMARY_KEYS, DBSCHEMA_FOREIGN_KEYS
USER_CONSTRAINTS	SQLPrimaryKeys, SQLForeignKeys	DBSCHEMA_PRIMARY_KEYS, DBSCHEMA_FOREIGN_KEYS
USER_IND_COLUMNS	SQLStatistics	DBSCHEMA_STATISTICS
USER_INDEXES	SQLStatistics	DBSCHEMA_STATISTICS
USER_OBJECTS	SQLTables, SQLProcedures, SQLStatistics	DBSCHEMA_TABLES, DBSCHEMA_PROCEDURES, DBSCHEMA_STATISTICS
USER_TAB_COLUMNS	SQLColumns	DBSCHEMA_COLUMNS
USER_TAB_COMMENTS	SQLTables	DBSCHEMA_TABLES
USER_TABLES	SQLStatistics	DBSCHEMA_STATISTICS
USER_USERS	SQLTables	DBSCHEMA_TABLES
USER_VIEWS	SQLTables	DBSCHEMA_TABLES

Generic Connectivity Data Dictionary Descriptions

The Generic Connectivity data dictionary tables and views provide this information:

- Name, datatype, 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 Oracle 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

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2(30)	-
TABLE_NAME	NOT NULL	VARCHAR2(30)	-
COLUMN_NAME	NOT NULL	VARCHAR2(30)	-
COMMENTS	-	VARCHAR2(4000)	NULL

ALL_CONS_COLUMNS

Name	Null?	Type	Value
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

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 "NO ACTION" 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

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	-

Name	Null?	Type	Value
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)	-
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

Name	Null?	Type	Value
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	-	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

Name	Null?	Type	Value
DOMIDX_STATUS	-	VARCHAR2 (12)	NULL
DOMIDX_OPSTATUS	-	VARCHAR2 (6)	NULL
FUNCIDX_STATUS	-	VARCHAR2 (8)	NULL

ALL_OBJECTS

Name	Null?	Type	Value
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
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

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2 (30)	-
TABLE_NAME	NOT NULL	VARCHAR2 (30)	-
COLUMN_NAME	NOT NULL	VARCHAR2 (30)	-
DATA_TYPE	-	VARCHAR2 (106)	-

Name	Null?	Type	Value
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
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

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2 (30)	-
TABLE_NAME	NOT NULL	VARCHAR2 (30)	-

Name	Null?	Type	Value
TABLE_TYPE	-	VARCHAR2(11)	"TABLE" or "VIEW"
COMMENTS	-	VARCHAR2(4000)	NULL

ALL_TABLES

Name	Null?	Type	Value
OWNER	NOT NULL	VARCHAR2(30)	-
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

Name	Null?	Type	Value
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
PARTITIONED	-	VARCHAR2(3)	NULL
IOT_TYPE	-	VARCHAR2(12)	NULL
TEMPORARY	-	VARCHAR2(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

Name	Null?	Type	Value
USERNAME	NOT NULL	VARCHAR2(30)	-
USER_ID	NOT NULL	NUMBER	0

Name	Null?	Type	Value
CREATED	NOT NULL	DATE	NULL

ALL_VIEWS

Name	Null?	Type	Value
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
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)	-
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

Name	Null?	Type	Value
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)	-
TABLE_NAME	-	VARCHAR2 (30)	-
COLUMN_NAME	-	VARCHAR2 (4000)	-
COLUMN_POSITION	-	NUMBER	-
COLUMN_LENGTH	-	NUMBER	-
DESCEND	-	VARCHAR2 (4)	"DESC " or "ASC "

USER_INDEXES

Name	Null?	Type	Value
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

Name	Null?	Type	Value
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
LAST_ANALYZED	-	DATE	NULL
DEGREE	-	VARCHAR2(40)	NULL
INSTANCES	-	VARCHAR2(40)	NULL
PARTITIONED	-	VARCHAR2(3)	NULL
TEMPORARY	-	VARCHAR2(1)	NULL
GENERATED	-	VARCHAR2(1)	NULL

Name	Null?	Type	Value
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
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

Name	Null?	Type	Value
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)	-
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 (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_DECL_LENGTH	-	NUMBER	0
GLOBAL_STATS	-	VARCHAR2 (3)	NULL

Name	Null?	Type	Value
USER_STATS	-	VARCHAR2 (3)	NULL
AVG_COL_LEN	-	NUMBER	0

USER_TAB_COMMENTS

Name	Null?	Type	Value
TABLE_NAME	NOT NULL	VARCHAR2 (30)	-
TABLE_TYPE	-	VARCHAR2 (11)	"TABLE" or "VIEW"
COMMENTS	-	VARCHAR2 (4000)	NULL

USER_TABLES

Name	Null?	Type	Value
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

Name	Null?	Type	Value
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
PARTITIONED	-	VARCHAR2 (3)	NULL
IOT_TYPE	-	VARCHAR2 (12)	NULL
TEMPORARY	-	VARHCAR2 (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

Name	Null?	Type	Value
USERNAME	NOT NULL	VARCHAR2 (30)	-
USER_ID	NOT NULL	NUMBER	0
ACCOUNT_STATUS	NOT NULL	VARCHAR2 (32)	OPEN
LOCK_DATE	-	DATE	NULL
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

Name	Null?	Type	Value
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

agent control utility. See `agctl`.

agents

- Generic Connectivity, 2-4
- Heterogeneous Services
 - architecture, 2-2
 - disabling self-registration, 4-14
 - registering, 4-10, 4-11, 4-12
 - types of agents, 2-3
- multithreaded Heterogeneous Services, 5-1
- specifying initialization parameters for, 4-3
- Transparent Gateways, 2-3

`agctl`, 5-3, 5-6

- commands, 5-7
- shell mode commands, 5-11
- single-line command mode, 5-8

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, 4-7
 - using OCI for bulk fetches, 4-8

array fetches, 4-8

- agents, 4-9

B

bind queries

executing using pass-through SQL, 3-11

`BIND_INOUT_VARIABLE` procedure, 3-7, 3-11

`BIND_OUT_VARIABLE` procedure, 3-6, 3-10

`BIND_VARIABLE` procedure, 3-6

buffers

- multiple rows, 3-12

bulk fetches

- optimizing data transfers using, 4-7

C

character sets

- Heterogeneous Services, A-8

`CLOSE_CURSOR` function, 3-6

commit point site

- commit point strength, A-3

configuring

- Generic Connectivity, 7-7
- Transparent Gateways, 4-2

copying data

- from Oracle database server to SQL Server, 4-18
- from SQL Server to Oracle database server, 4-19

`INSERT` statement, 4-19

D

data dictionary

- contents with Generic Connectivity, C-5
- mapping for Generic Connectivity, C-5
- Oracle server name/SQL Server name, C-5
- translation support for Generic Connectivity, C-1

data dictionary views

- Generic Connectivity, C-5

- Heterogeneous Services, 4-20, C-2
- database links
 - heterogeneous systems, 4-3
- date formats
 - Heterogeneous Services, A-9, A-11
- DBMS_HS_PASSTHROUGH package, 3-5
 - list of functions and procedures, 3-6
- describe cache high water mark
 - definition, A-5
- dispatcher threads
 - multithreaded Heterogeneous Services agents, 5-3, 5-5
- distributed queries
 - optimizing performance, 6-2
- drivers
 - ODBC, 7-12
 - OLE DB (FS), 7-15
 - OLE DB (SQL), 7-14
- dynamic performance views
 - Heterogeneous Services, 4-25
 - determining open sessions, 4-26
 - determining which agents are on host, 4-26

E

- EXECUTE_IMMEDIATE procedure, 3-7
 - restrictions, 3-7
- EXECUTE_NON_QUERY procedure, 3-7

F

- FDS_CLASS, 4-12
- FDS_CLASS_VERSION, 4-12
- FDS_INST_NAME, 4-13
- FETCH_ROW procedure, 3-7
 - executing queries using pass-through SQL, 3-11
- fetches
 - bulk, 4-7
 - optimizing round-trips, 3-12

G

- gateways
 - how they work, 2-9
- Generic Connectivity

- architecture, 7-2
 - Oracle and non-Oracle on same machine, 7-4
 - Oracle and non-Oracle on separate machines, 7-3
- configuration, 7-7
- creating initialization file, 7-7
- data dictionary
 - translation support, C-1
- defined, 1-3
- definition, 7-2
- DELETE statement, 7-6
- editing initialization file, 7-7
- Heterogeneous Services, 2-4
- INSERT statement, 7-6
- non-Oracle data dictionary access, C-2
- ODBC connectivity requirements, 7-12
- OLE DB (FS) connectivity requirements, 7-15
- OLE DB (SQL) connectivity requirements, 7-14
- restrictions, 7-5
- setting parameters for ODBC source, 7-9
 - UNIX, 7-10
 - Windows NT, 7-9
- setting parameters for OLE DB source, 7-11
- SQL execution, 7-5
- supported functions, 7-6
- supported SQL syntax, 7-6
- types of agents, 7-2
- UPDATE statement, 7-6

generic connectivity

- error tracing, A-7
- GET_VALUE procedure, 3-7, 3-10, 3-11
- globalization support
 - Heterogeneous Services, A-7
 - character set of non-Oracle source, A-10
 - date format, A-9
 - languages in character date values, A-10

H

- heterogeneous distributed systems
 - accessing, 4-2
- Heterogeneous Services
 - agent control utility (agtctl), 5-7
 - agent registration, 4-10
 - avoiding configuration mismatches, 4-11

- disabling, 4-14
- enabling, 4-10
- self-registration, 4-12
- 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, 4-7
 - using OCI for bulk fetches, 4-8
- creating database links, 4-3
- data dictionary views, 4-20, C-2
 - types, 4-20
 - understanding sources, 4-21
 - using general views, 4-22
 - using SQL service views, 4-24
 - using transaction service views, 4-23
- defining maximum number of open cursors, A-12
- dynamic performance views, 4-25
 - V\$HS_AGENT view, 4-26
 - V\$HS_SESSION view, 4-26
- Generic Connectivity
 - architecture, 7-2
 - creating initialization file, 7-7
 - definition, 7-2
 - editing initialization file, 7-7
 - non-Oracle data dictionary access, C-2
 - ODBC connectivity requirements, 7-12
 - OLE DB (FS) connectivity requirements, 7-15
 - OLE DB (SQL) connectivity requirements, 7-14
 - restrictions, 7-5
 - setting parameters for ODBC source, 7-9
 - setting parameters for OLE DB source, 7-11
 - SQL execution, 7-5
 - supported functions, 7-6
 - supported SQL syntax, 7-6
 - supported tables, C-5
 - types of agents, 7-2
- initialization parameters, 2-6, 4-5, 7-7, A-1
- multithreaded agents, 5-1
- optimizing data transfer, A-12
- setting global name, A-5
- setting up access using Transparent Gateway, 4-2
- setting up environment, 4-2
- specifying cache high water mark, A-5
- specifying cache size, A-12
- specifying commit point strength, A-3
- SQL service, 2-5
- testing connections, 4-4
- transaction service, 2-4
- tuning internal data buffering, A-13
- tuning LONG data transfer, A-9
- HS_AUTOREGISTER initialization parameter
 - using to enable agent self-registration, 4-14
- HS_BASE_CAPS view, 4-20
- HS_BASE_DD view, 4-20
- HS_CALL_NAME initialization parameter, A-3
- HS_CLASS_CAPS view, 4-21
- HS_CLASS_DD view, 4-21
- HS_CLASS_INIT view, 4-21
- HS_COMMIT_POINT_STRENGTH initialization parameter, A-3
- HS_DB_NAME initialization parameter, A-5
- HS_DESCRIBE_CACHE_HWM initialization parameter, A-5
- HS_FDS_CLASS view, 4-21
- HS_FDS_CONNECT_INFO initialization parameter, A-5
 - specifying connection information, 7-8
- HS_FDS_DEFAULT_SCHEMA_NAME initialization parameter, A-6
- HS_FDS_FETCH_ROWS initialization parameter, 4-9
- HS_FDS_INST view, 4-21
- HS_FDS_SHAREABLE_NAME initialization parameter, A-7
- HS_FDS_TRACE_LEVEL initialization parameter, A-7
 - enabling agent tracing, 7-8
- HS_LANGUAGE initialization parameter, A-7
- HS_LONG_PIECE_TRANSFER_SIZE initialization parameter, A-9
- HS-NLS_DATE_FORMAT initialization parameter, A-9

- HS_NLS_DATE_LANGUAGE initialization parameter, A-10
- HS_NLS_NCHAR initialization parameter, A-10
- HS_NLS_NUMERIC_CHARACTER initialization parameter, A-10
- HS_NLS_TIMESTAMP_FORMAT initialization parameter, A-11
- HS_NLS_TIMESTAMP_TZ_FORMAT initialization parameter, A-11
- HS_OPEN_CURSORS initialization parameter, A-12
- HS_ROWID_CACHE_SIZE initialization parameter, A-12
- HS_RPC_FETCH_REBLOCKING initialization parameter, 4-9, A-12
- HS_RPC_FETCH_SIZE initialization parameter, 4-9, A-13
- HS_TIME_ZONE initialization parameter, A-14

I

- IFILE initialization parameter, A-14
- information integration
 - benefits of Oracle solutions, 1-4
 - challenges, 1-2
 - Generic Connectivity, 1-3
 - how Oracle addresses, 1-2
 - Messaging Gateway, 1-4
 - Open System Interfaces, 1-4
 - Oracle Streams, 1-4
 - Oracle Transparent Gateways, 1-3
- initialization parameters
 - Heterogeneous Services (HS), 2-6, 4-5, 7-7, A-1
- initialization parameters (HS)
 - common to all gateways, 4-5
 - descriptions, A-1
 - Generic Connectivity, 7-7
 - purpose, 2-6

L

- listeners, 4-2

M

- Messaging Gateway
 - defined, 1-4
- monitor thread
 - multithreaded Heterogeneous Services agents, 5-3, 5-5
- multiple rows
 - buffering, 3-12
- multithreaded Heterogeneous Services agents
 - administering, 5-6
 - advantages, 5-2
 - agent control utility (agtctl), 5-7
 - architecture, 5-3
 - configuration parameters, 5-12
 - dispatcher threads, 5-3, 5-5
 - monitor thread, 5-3, 5-5
 - task threads, 5-3, 5-6

O

- OCI
 - optimizing data transfers using, 4-8
- ODBC agents
 - connectivity requirements, 7-12
 - functions, 7-12
- ODBC connectivity
 - data dictionary mapping, C-5
 - ODBC driver, 7-12
 - requirements, 7-12
 - specifying connection information
 - UNIX, A-6
 - Windows NT, A-6
 - specifying path to library, A-7
- OLE DB (FS) drivers, 7-15
- OLE DB (SQL) drivers, 7-14
- OLE DB agents
 - connectivity requirements, 7-14, 7-15
- OLE DB connectivity
 - data dictionary mapping, C-5
 - setting connection information, A-6
- OLE DB drivers
 - data provider requirements, 7-15
 - initialization properties, 7-17
 - rowset properties, 7-17

Open System Interfaces

defined, 1-4

OPEN_CURSOR procedure, 3-6

Oracle database server

SQL construct processing, 4-14

Oracle Net Services listener, 2-3, 4-2

Oracle precompiler

optimizing data transfers using, 4-8

Oracle Streams

defined, 1-4

Oracle Transparent Gateways

defined, 1-3

optimizing SQL statements, 6-2

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, 3-11

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

remote functions

referenced in SQL statements, A-3

rows

buffering multiple, 3-12

S

SELECT statement

accessing non-Oracle system, C-2

service names

specifying in database links, 4-3

SQL capabilities

data dictionary tables, 4-24

SQL service

data dictionary views, 2-8, 4-20

Heterogeneous Services, 2-5

views

Heterogeneous Services, 4-24

SQL statements

optimizing distributed, 6-2

Streams

using for heterogeneous connectivity, 3-3

Synonyms, 4-17

T

task threads

multithreaded Heterogeneous Services

agents, 5-3, 5-6

transaction service

Heterogeneous Services, 2-4

views

Heterogeneous Services, 4-23

transparent gateways

accessing Heterogeneous Services agents, 4-2

creating database links, 4-3

testing connections, 4-4

U

unsupported functions

Generic Connectivity, 7-6

user-defined functions

executing on non-Oracle database, 4-15

V

V\$HS_AGENT view

determining which agents are on host, 4-26

V\$HS_PARAMETER view

listing HS parameters, 4-27

V\$HS_SESSION view
 determining open sessions, 4-26
variables
 bind, 3-7