

Accessing Databases

iPlanet™ Unified Development Server

Version 5.0

August 2001

Copyright (c) 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Sun, Sun Microsystems, the Sun logo, Forte, iPlanet, Unified Development Server, and the iPlanet logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Federal Acquisitions: Commercial Software - Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright (c) 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats - Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Sun, Sun Microsystems, le logo Sun, Forte, iPlanet, Unified Development Server, et le logo iPlanet sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Contents

List of Figures	11
List of Tables	13
List of Procedures	15
Preface	17
Product Name Change	17
Audience for This Guide	18
Organization of This Guide	18
Text Conventions	19
Other Documentation Resources	20
iPlanet UDS Documentation	20
Express Documentation	21
WebEnterprise and WebEnterprise Designer Documentation	21
Online Help	21
iPlanet UDS Example Programs	21
Viewing and Searching PDF Files	22
Chapter 1 Introduction to the iPlanet UDS Database Interface	25
About the iPlanet UDS Database Interface	26
Support for Multiple Database Management Systems	26
Support for Standard and Proprietary SQL	27
ANSI SQL and Vendor SQL Extensions	27
Accessing Databases Using TOOL SQL	28
Accessing Databases using the Database Classes	28
The iPlanet UDS GenericDBMS Library	29
How to Access an RDBMS	30

Database Vendor Notes	32
DB2 Notes	32
Informix Notes	33
Oracle Notes	33
Rdb Notes	34
Unsupported Database Features	35
Using ODBC	35
Supported Data Sources	36
Supported Database Drivers	37
Using ODBC When Running iPlanet UDS in Standalone Model	37
Cautions on Using ODBC with iPlanet UDS	38
Chapter 2 Defining a Resource Manager	39
About Resource Managers	39
Choosing a Node for the Resource Name	41
Defining a Resource Name	42
Using Escript to Define a Resource Name	44
Environment Variables	44
DB2 Variables	46
Informix Variables	46
Oracle Variables	46
Rdb Variables	47
Sybase Variables	47
Removing a Resource Name	47
Testing a Resource Name with the DynamicSQL Example	48
Using a Local Database when Running iPlanet UDS Standalone	52
Setting up an iPlanet UDS ODBC Data Source on Windows	53
Using an Existing Data Source	54
Chapter 3 Making a Database Connection	55
Connecting to a Database	55
Using Database Service Objects	56
When to Use a DBSession Service Object	57
When to Use a DBResourceMgr Service Object	57
When You Should Not Use a Service Object	58
Creating a Database Service Object	58
Entering General Properties	59
Entering Database Information	60
Entering Search Path Information	62
Entering Connection Information	62
Specifying Connection Options in a Method	63
Optimizing Service Object Performance	63
Visibility of the Service Object	65

Making a Database Connection	65
Connecting with a DBSession Service Object	65
Connecting with a DBResourceMgr Service Object	66
Connecting to a Database without a Service Object	68
Dynamically Choosing a Database Vendor	69
Other Connection Information	71
Using Variable User Names and Passwords	71
Creating a DefaultDBSession Service Object	73
Reconnecting to a Database Session	75
Disconnecting a Database Session	76
Vendor-Specific Notes	76
Informix	76
Oracle	77
Rdb	78
Chapter 4 Working with Data Types	79
Using Database Data with iPlanet UDS	79
Using Simple Data Types in TOOL	80
Using Nullable DataValue Subclasses	80
Data Type Conversion	81
Reading the Data Type Conversion Tables	82
DB2 Data Conversion Table and Notes	84
Informix Data Conversion Table and Notes	85
ODBC Data Conversion Table and Notes	86
Oracle Data Conversion Table and Notes	88
Rdb Data Conversion Table and Notes	89
Sybase Data Conversion Table and Notes	92
Chapter 5 Manipulating Data	93
Accessing Database Data from iPlanet UDS	93
Using iPlanet UDS Names in SQL Statements	95
TOOL SQL Statements	95
Using Conditional TOOL for Vendor-Specific Code	96
Using TOOL Statements to Query Data	98
Selecting Data and Object Creation	98
Selecting a Single Row	100
Selecting into a Variable	100
Selecting into an Object	100
When Attribute and Column Names Match	101
When Attribute and Column Names Do Not Match	101
When Attributes are of Other Class Types	102
When Attributes are Inherited	104

Using TOOL Statements to Query Data <i>(continued)</i>	
Selecting Multiple Rows into Arrays	104
Selecting Multiple Rows Using the TOOL for Statement	105
Selecting Multiple Rows using Cursors	105
Defining a Cursor	106
Retrieving Rows	107
Fetching into an Array	108
Fetching an Arbitrary Number of Rows	108
Repeating a Statement Block	108
Using TOOL to Update Data	109
Inserting a Single Row	109
Inserting Variables	109
Inserting from an Object	109
Inserting Multiple Rows	110
Updating a Row	110
Deleting a Row	111
Executing a Single SQL Statement	111
Executing a Database Procedure	111
Vendor-Specific Notes on Database Procedures	112
Working with ImageData Objects	113
Using Binary Large Objects (BLOBs)	113
Selecting Binary Data	113
Inserting Binary Data	114
Vendor-Specific Notes on BLOB Handling	116
Using iPlanet UDS Classes to Execute SQL	118
DBSession Methods	118
Executing Single SQL Statements	119
Using Prepared Statements	120
Executing Prepared Queries	120
About the DynamicDataAccess Example	121
Building the SQL Statement	122
Preparing the Statement	124
Opening the Cursor	125
Fetching Rows from the Result Set	126
Storing the Data	126
Closing the Cursor	127
Executing Prepared DML Statements	127
About the DynamicDataAccess Example	127
Building the SQL Statement	128
Preparing the Statement	129
Processing Placeholders	130
Executing the Statement	131
Removing the Statement	131

Improving Application Performance	132
Multi-threaded Database Access	132
Enabling Single-Threaded Access to Databases	133
Mapping DBDataSets into Objects	133
Vendor-Specific Information	136
Informix	136
Scroll Cursor Support	136
Chapter 6 Transactions	139
Relationship Between iPlanet UDS and Database Transactions	139
Explicit iPlanet UDS Transactions	140
Implicit iPlanet UDS Transactions	141
Single SQL DML Statements	142
SQL Execute Immediate Statements with DDL	142
Use of Cursors in Implicit iPlanet UDS Transactions	143
SQL Open Cursor	143
SQL Fetch Cursor	144
For Loops	145
SQL Close Cursor	145
Independent, Dependent, and Nested Transactions	145
Using Dependent Transactions	145
Using Independent Transactions	146
Avoid Nested Transactions	146
Common Problems with Shared and Transactional Objects	147
Unexpected Blocking Due to a Long-Running Query	148
Unexpected Blocking Due to a Long-Running Transaction	148
Avoiding Deadlocks	149
Transactions and Database Sessions	150
Multitasking in a Database Session	150
iPlanet UDS Support for Two-Phase Commit	151
Two-Phase Commit with One Database Vendor	153
Two-Phase Commit with Multiple Database Vendors	153
Notes on Vendor-Specific Transaction Handling	154
DB2	154
Informix	154
Rdb	155
Chapter 7 Error Handling	157
Types of Database Exceptions	158

Appendix A Database Example Applications	161
How to Install iPlanet UDS Example Applications	161
Overview of Database Example Applications	162
GenericDBMS Library Examples	162
Application Descriptions	163
DynamicDataAccess	163
DynamicSQL	164
WinDB	165
Appendix B TOOL SQL Statement Reference	167
Note on Vendor-Specific SQL Extensions	168
SQL Close Cursor	168
Syntax	168
Example	168
Description	168
SQL Delete	169
Syntax	169
Example	169
Description	169
Return Value	169
Table Name	170
Where Clause	170
On Session Clause	170
SQL Execute Immediate	171
Syntax	171
Example	171
Description	171
Return Value	172
On Session Clause	173
SQL Execute Procedure	173
Syntax	173
Example	173
Description	174
Return Value	174
Procedure Name	174
Parameter List	175
On Session Clause	176
Exceptions	176

SQL Fetch Cursor	176
Syntax	176
Example	176
Description	176
Return Value	178
Cursor	179
Into Clause	179
SQL Insert	180
Syntax	180
Example	180
Description	180
Return Value	181
Table Name	181
The Column List	181
Specifying the Insert Values	181
On Session Clause	182
SQL Open Cursor	182
Syntax	183
Example	183
Description	183
Cursor Reference	184
Placeholder Assignment	184
On Session Clause	185
SQL Select	185
Syntax	185
Example	186
Description	186
Return Value	186
Eliminating Duplicate Rows	187
Column List	187
Into Clause	188
From Clause	189
Where Clause	189
Group By Clause	189
Having Clause	189
Order By Clause	189
On Session Clause	190
Exceptions	190

SQL Update	190
Syntax	190
Example	190
Description	191
Return Value	191
Table Name	191
Set Clause	191
Where Clause	191
On Session Clause	192
Index	193

List of Figures

Figure 1-1	Components of a Typical iPlanet UDS Database Application	30
Figure 1-2	Steps to Build an iPlanet UDS Database Application	31
Figure 1-3	ODBC Architecture.	36
Figure 2-1	Resource Name used by Two Databases	40
Figure 3-1	New Service Object Dialog	59
Figure 3-2	General Properties Tab Page	59
Figure 3-3	Database Tab Page	60
Figure 3-4	Search Path Tab Page	62
Figure 3-5	Connection Tab Page	63
Figure 3-6	Partitioning with Intermediate Data Service on Same Partition as Session Service	64
Figure 3-7	Login screen for DynamicData Access Example	68
Figure 3-8	Using Environment Variables in Service Object Definition	72
Figure 5-1	Choosing Selection Criteria	121
Figure 5-2	Dynamically Constructed SQL Statement	122
Figure 5-3	Displaying the Results	122
Figure 5-4	Insert Screen	128
Figure 7-1	The Database Exception Hierarchy	158

List of Tables

Table 1-1	Equivalent TOOL SQL Statements and Methods	27
Table 4-1	Legend for the Data Type Conversion Tables	82
Table 4-2	DB2 Data Conversion Table	84
Table 4-3	Informix Data Conversion Table	86
Table 4-4	ODBC Data Conversion Table	87
Table 4-5	Oracle Data Conversion Table	88
Table 4-6	Rdb Data Conversion Table	90
Table 4-7	Sybase Data Conversion Table	92

List of Procedures

- To copy the documentation to a client or server 22
- To view and search the documentation 22
- To define a resource name for a database using Environment Console 42
- To test a resource name 48
- To define a new ODBC data source 53
- To make the data source available to iPlanet UDS users 54
- To start a database session using a DBSession service object 66
- To start a database session using a DBResourceMgr service object 67
- To create and connect through a DBSession object 69
- To create a default database session for an entire project 73
- To create a default database session for a class and its associated methods 74
- To create a default database for use in a single method 75
- To make a connection using SQL*Net V2 77
- To execute a prepared select statement 120
- To execute a prepared insert, update, or delete statement 127
- To maximize your application's performance, remember to 132
- To use DynamicDataAccess 164
- To use DynamicSQL 164
- To use WinDB 166

Preface

Accessing Databases provides information about accessing databases from iPlanet™ Unified Development Server (iPlanet UDS) applications. This preface contains the following sections:

- “Product Name Change” on page 17
- “Audience for This Guide” on page 18
- “Organization of This Guide” on page 18
- “Text Conventions” on page 19
- “Other Documentation Resources” on page 20
- “iPlanet UDS Example Programs” on page 21
- “Viewing and Searching PDF Files” on page 22

Product Name Change

Forte 4GL has been renamed the iPlanet Unified Development Server. You will see full references to this name, as well as the abbreviations iPlanet UDS and UDS.

Audience for This Guide

System managers will find instructions for preparing a database management system to work with an iPlanet UDS application. This primarily entails defining a resource manager for each database that will interact with an iPlanet UDS program (described in [Chapter 2, “Defining a Resource Manager”](#)).

iPlanet UDS *application programmers* will find information they need to write applications using TOOL SQL statements or dynamic SQL methods.

We assume that these readers:

- have programming experience
- are familiar with their particular window system
- are familiar with SQL and their particular database management system
- understand the basic concepts of object-oriented programming as described in *A Guide to the iPlanet UDS Workshops*
- have used the iPlanet UDS workshops to create classes

Organization of This Guide

The following table briefly describes the contents of each chapter:

Chapter	Description
Chapter 1, “Introduction to the iPlanet UDS Database Interface”	Provides an overview of iPlanet UDS support for database integration.
Chapter 2, “Defining a Resource Manager”	Describes how to create a resource manager for database access.
Chapter 3, “Making a Database Connection”	Describes DBSession and DBResourceMgr services objects, how to create each, and their advantages and disadvantages.
Chapter 4, “Working with Data Types”	Discusses the mappings between database data types and iPlanet UDS TOOL data types, and points out conversion issues.
Chapter 5, “Manipulating Data”	Provides information about writing TOOL SQL statements and SQL methods to access (query) and manipulate (update) database data.

Chapter	Description
Chapter 6, “Transactions”	Discusses the relationship between database transactions and iPlanet UDS transactions.
Chapter 7, “Error Handling”	Provides a brief overview of iPlanet UDS’s database exception classes.
Appendix A, “Database Example Applications”	Describes the standard iPlanet UDS sample applications for database access, with instructions for installation and use.
Appendix B, “TOOL SQL Statement Reference”	Provides the primary source of information about the SQL subset of TOOL statements; for other TOOL statement reference, see the <i>TOOL Reference Guide</i> .

Text Conventions

This section provides information about the conventions used in this document.

Format	Description
<i>italics</i>	Italicized text is used to designate a document title, for emphasis, or for a word or phrase being introduced.
monospace	Monospace text represents example code, commands that you enter on the command line, directory, file, or path names, error message text, class names, method names (including all elements in the signature), package names, reserved words, and URLs.
ALL CAPS	Text in all capitals represents environment variables (FORTE_ROOT) or acronyms (UDS, JSP, iMQ). Uppercase text can also represent a constant. Type uppercase text exactly as shown.
Key+Key	Simultaneous keystrokes are joined with a plus sign: Ctrl+A means press both keys simultaneously.
Key-Key	Consecutive keystrokes are joined with a hyphen: Esc-S means press the Esc key, release it, then press the S key.

Other Documentation Resources

In addition to this guide, there are additional documentation resources, which are listed in the following sections. The documentation for all iPlanet UDS products (including Express, WebEnterprise, and WebEnterprise Designer) can be found on the iPlanet UDS Documentation CD. Be sure to read “[Viewing and Searching PDF Files](#)” on page 22 to learn how to view and search the documentation on the iPlanet UDS Documentation CD.

iPlanet UDS documentation can also be found online at <http://docs.iplanet.com/docs/manuals/uds.html>.

The titles of the iPlanet UDS documentation are listed in the following sections.

iPlanet UDS Documentation

- *A Guide to the iPlanet UDS Workshops*
- *Accessing Databases*
- *Building International Applications*
- *Esript and System Agent Reference Guide*
- *Fscript Reference Guide*
- *Getting Started With iPlanet UDS*
- *Integrating with External Systems*
- *iPlanet UDS Java Interoperability Guide*
- *iPlanet UDS Programming Guide*
- *iPlanet UDS System Installation Guide*
- *iPlanet UDS System Management Guide*
- *Programming with System Agents*
- *TOOL Reference Guide*
- *Using iPlanet UDS for OS/390*

Express Documentation

- *A Guide to Express*
- *Customizing Express Applications*
- *Express Installation Guide*

WebEnterprise and WebEnterprise Designer Documentation

- *A Guide to WebEnterprise*
- *Customizing WebEnterprise Designer Applications*
- *Getting Started with WebEnterprise Designer*
- *WebEnterprise Installation Guide*

Online Help

When you are using an iPlanet UDS development application, press the F1 key or use the Help menu to display online help. The help files are also available at the following location in your iPlanet UDS distribution:

```
FORTE_ROOT/userapp/forte/cln/*.hlp.
```

When you are using a script utility, such as Fscript or Escript, type help from the script shell for a description of all commands, or help *<command>* for help on a specific command.

iPlanet UDS Example Programs

A set of example programs is shipped with the iPlanet UDS product. The examples are located in subdirectories under `$FORTE_ROOT/install/examples`. The files containing the examples have a `.pex` suffix. You can search for TOOL commands or anything of special interest with operating system commands. The `.pex` files are text files, so it is safe to edit them, though you should only change private copies of the files.

Viewing and Searching PDF Files

You can view and search iPlanet UDS documentation PDF files directly from the documentation CD-ROM, store them locally on your computer, or store them on a server for multiuser network access.

NOTE You need Acrobat Reader 4.0+ to view and print the files. Acrobat Reader with Search is recommended and is available as a free download from <http://www.adobe.com>. If you do not use Acrobat Reader with Search, you can only view and print files; you cannot search across the collection of files.

➤ **To copy the documentation to a client or server**

1. Copy the `doc` directory and its contents from the CD-ROM to the client or server hard disk.

You can specify any convenient location for the `doc` directory; the location is not dependent on the iPlanet UDS distribution.

2. Set up a directory structure that keeps the `udsdoc.pdf` and the `uds` directory in the same relative location.

The directory structure must be preserved to use the Acrobat search feature.

NOTE To uninstall the documentation, delete the `doc` directory.

➤ **To view and search the documentation**

1. Open the file `udsdoc.pdf`, located in the `doc` directory.
2. Click the Search button at the bottom of the page or select Edit > Search > Query.

3. Enter the word or text string you are looking for in the Find Results Containing Text field of the Adobe Acrobat Search dialog box, and click Search.

A Search Results window displays the documents that contain the desired text. If more than one document from the collection contains the desired text, they are ranked for relevancy.

NOTE For details on how to expand or limit a search query using wild-card characters and operators, see the Adobe Acrobat Help.

4. Click the document title with the highest relevance (usually the first one in the list or with a solid-filled icon) to display the document.

All occurrences of the word or phrase on a page are highlighted.

5. Click the buttons on the Acrobat Reader toolbar or use shortcut keys to navigate through the search results, as shown in the following table:

Toolbar Button	Keyboard Command
Next Highlight	Ctrl+]]
Previous Highlight	Ctrl+[[
Next Document	Ctrl+Shift+]]

To return to the `udsdoc.pdf` file, click the Homepage bookmark at the top of the bookmarks list.

6. To revisit the query results, click the Results button at the bottom of the `udsdoc.pdf` home page or select Edit > Search > Results.

Introduction to the iPlanet UDS Database Interface

This chapter provides an introduction to the iPlanet UDS database interface. Database management systems (DBMS) are standard components of many complex applications. iPlanet UDS includes interfaces to the following database management systems:

- IBM UDB/DB2
- Informix Dynamic Server
- Microsoft SQL Server
- Oracle
- Rdb (both Oracle Rdb and its predecessor, DEC Rdb)
- Sybase

In addition, the iPlanet UDS database interface also supports Microsoft's Open Database Connectivity (ODBC), allowing access to data sources through database drivers that conform to the ODBC standard.

For information about which database versions or platforms are certified for a given iPlanet UDS release, you should refer to the platform support matrix, available at <http://www.forte.com/support/platforms.html>.

About the iPlanet UDS Database Interface

The iPlanet UDS database interface allows iPlanet UDS application developers to build database access into their applications quickly and easily. Application developers can take advantage of virtually all database features, both standard and proprietary, while building applications that are portable across hardware platforms, operating systems, and databases.

The primary iPlanet UDS tools that support database access are the TOOL SQL statements and iPlanet UDS's GenericDBMS library. This introduction describes these tools in more detail. iPlanet UDS also provides a number of sample applications that demonstrate working with database data. These programs are described in [Appendix A, "Database Example Applications."](#)

Support for Multiple Database Management Systems

The iPlanet UDS database interface supports the following database systems:

- IBM UDB/DB2 R6000 & Mainframe
- Informix
- Microsoft SQL Server
- ODBC Level 1 API ODBC 2.1 SDK
- Oracle
- Rdb
- Sybase

For detailed information regarding which versions and platforms of a particular database management system are supported, refer to the *iPlanet UDS System Installation Guide*, which contains the following types of information:

- version and release numbers of certified third party components
- notes on platforms that are certified for a given product and release

The *iPlanet UDS Release Notes* may also contain additional information.

When this manual describes information that is unique to one or more database vendors, the information is identified as *database vendor-specific* information.

Support for Standard and Proprietary SQL

iPlanet UDS database applications can contain any SQL statement supported by any database vendor, with very few exceptions. An application can contain any combination of DML (data manipulation language, including queries, updates, inserts, and deletes), DDL (data definition language, for creating, altering, and dropping tables), or DCL (data control language, for database access and security). Many applications consist primarily of DML.

ANSI SQL and Vendor SQL Extensions

You can write portable applications using ANSI standard SQL statements that all vendors support. You can also use vendor-specific SQL extensions, such as outer-joins and advanced set operations, although when you use such extensions, your applications may not be portable across all databases.

The iPlanet UDS database interface supports SQL access to a relational database system using two different approaches:

- using iPlanet UDS TOOL SQL statements
- using classes in iPlanet UDS's GenericDBMS library

Table 1-1 lists the SQL statements supported by TOOL along with the equivalent method. All methods are on the DBSession class of the GenericDBMS library.

Table 1-1 Equivalent TOOL SQL Statements and Methods

TOOL SQL Statement	DBSession Method	Description
sql select	Select *	Retrieve a single row from a database table or retrieves multiple rows from a database table into an array.
sql delete	Execute *	Remove rows from a database table.
sql insert	Execute *	Add a new row to a database table.
sql update	Execute *	Replace values in a database table.
sql open cursor	OpenCursor *	Open a cursor.
sql close cursor	CloseCursor *	Close a cursor.
sql fetch cursor	FetchCursor *	Retrieve a single row from a cursor or retrieve multiple rows from a cursor into an array.

* These statements must be previously prepared using the DBSession class interface.

** This statement must be previously prepared using the PrepareProcedure method of the DBSession class.

Table 1-1 Equivalent TOOL SQL Statements and Methods (*Continued*)

TOOL SQL Statement	DBSession Method	Description
sql execute immediate	ExecuteImmediate	Execute a single SQL statement specified as a literal string, a string variable, or a TextData variable.
sql execute procedure ** (all databases except Sybase)	OpenCursor (Sybase procedures only)	Execute a database procedure.

* These statements must be previously prepared using the DBSession class interface.
** This statement must be previously prepared using the PrepareProcedure method of the DBSession class.

Accessing Databases Using TOOL SQL

TOOL SQL statements can be used in a number of situations. If you have a choice of using a TOOL SQL statement or a GenericDBMS class, it is usually simpler to use the TOOL statement.

In general, any operation you can perform using TOOL SQL you can also perform using a GenericDBMS class. For single return values, use the TOOL SQL statement `execute procedure`; for multiple result sets, use the DBSession class.

[Appendix B, "TOOL SQL Statement Reference"](#) is the primary reference for TOOL SQL statements. Many of these statements are demonstrated in [Chapter 5, "Manipulating Data."](#) For reference information regarding non-SQL TOOL statements, refer to the *TOOL Reference Guide*.

Accessing Databases using the Database Classes

The GenericDBMS library contains a number of classes that support connecting to and working with relational database systems. This library must be named as a supplier plan in order to access any database. You use the Project Workshop to make the GenericDBMS library a supplier plan to the project in which you want to access the database.

Many GenericDBMS classes and methods, in particular, those on the DBSession class, are functionally equivalent to TOOL SQL statements. For example, when you are writing a method, you can either invoke the Select method on the DBSession class or use the TOOL statement **sql select**; either way provides the same execution, features, and performance.

Other classes in GenericDBMS, such as DBDatabaseDesc, DBTableDesc, DBColumnDesc, and DBKeyDesc are “helper” classes; they do not have TOOL SQL equivalents, but are very useful for returning information about database tables, columns, keys, and so on.

NOTE In some circumstances you must use the classes and methods in GenericDBMS. You must use the classes in the GenericDBMS library to build applications that execute SQL statements that are constructed at runtime. These statements are often used in applications to construct queries or updates based upon input from the end-user. TOOL SQL statements cannot be used to execute SQL statements specified at runtime.

“Using iPlanet UDS Classes to Execute SQL” on page 118 describes how to use the GenericDBMS classes to build and execute SQL statements at runtime.

The iPlanet UDS GenericDBMS Library

Every iPlanet UDS application that interacts with a database system must include the iPlanet UDS library called GenericDBMS as a supplier plan. This library contains the classes that allow you to use SQL in your applications.

In addition to the GenericDBMS library you also need a vendor-specific library for each database to which you actually connect. For example, if you partition your application to use an Oracle resource manager, then you need the library called Oracle. The database-specific libraries have the name of the database vendor (as in DB2, Informix, and so on).

Always add the GenericDBMS library as a supplier library for any database access, even if you know that you will only use one type of database. Do not explicitly refer to or add the individual database vendor libraries as supplier libraries. They are added dynamically during partitioning, as needed.

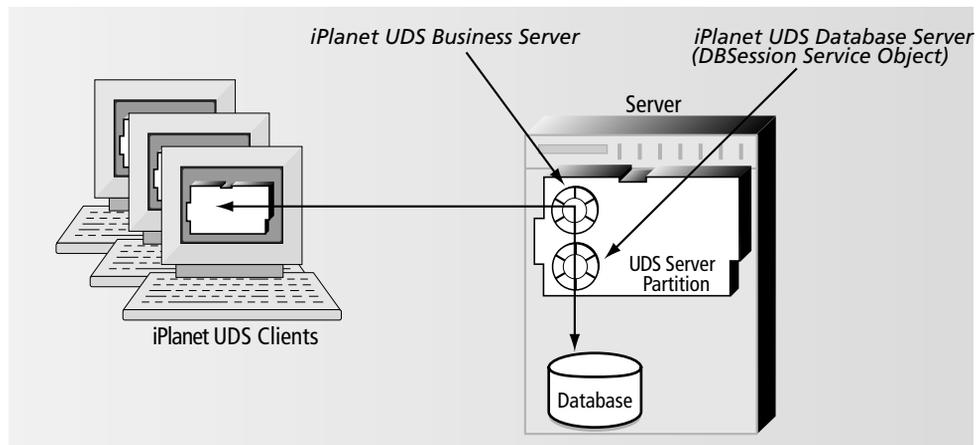
The iPlanet UDS online Help contains reference material for the GenericDBMS library.

How to Access an RDBMS

This manual describes how to integrate an iPlanet UDS application with a database, so that your iPlanet UDS application users can query and update database data. To integrate a database with an iPlanet UDS application, you can define a service object whose purpose is to handle calls to the database (you can also build iPlanet UDS applications without such a service object).

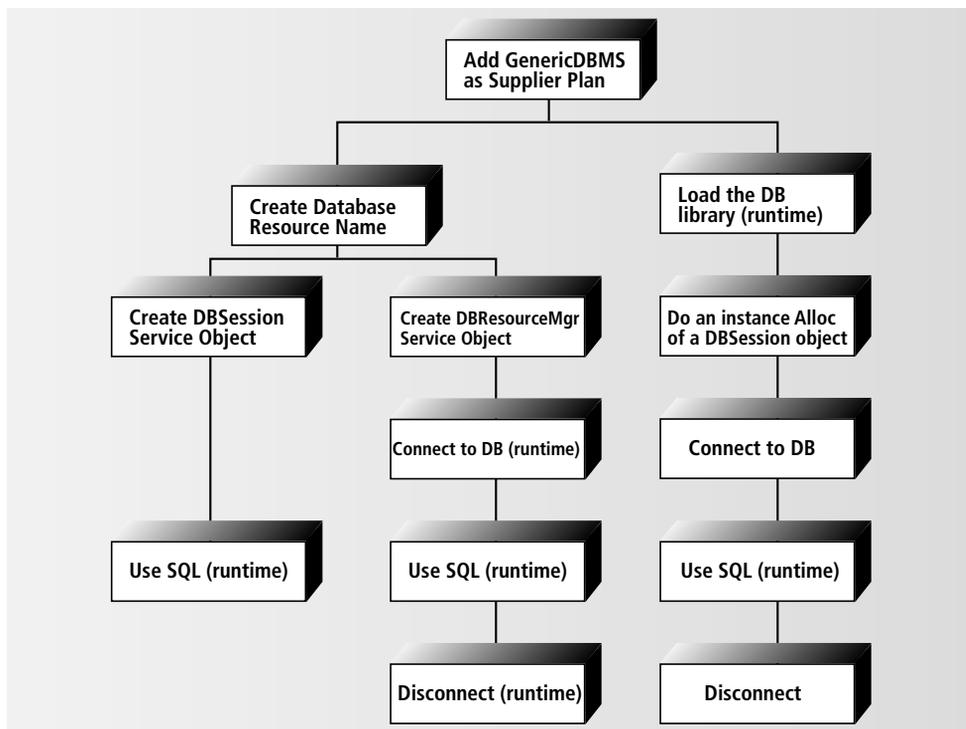
The basic components of an iPlanet UDS database application configuration are shown in [Figure 1-1](#).

Figure 1-1 Components of a Typical iPlanet UDS Database Application



[Figure 1-1](#) does not necessarily represent a typical iPlanet UDS database application, but rather the minimal components of one. Note that it shows two service objects. Each iPlanet UDS application service object (there may be several) manages iPlanet UDS objects and their logic. A DBSession service object (or a DBResourceMgr service object) coordinates communications with the database, and interacts with the iPlanet UDS service objects. Service objects are sometimes called “servers.”

[Figure 1-2](#) shows a chart of the high-level steps required to build an application such as that shown in [Figure 1-1](#). This book describes each of these tasks in more detail.

Figure 1-2 Steps to Build an iPlanet UDS Database Application

The following table lists some of the integration steps, keyed to sections in this manual.

Task	Reference
Include the GenericDBMS library as a supplier plan.	<i>A Guide to the iPlanet UDS Workshops</i>
Define a database resource name.	"Defining a Resource Name" on page 42
Set database environment variables.	"Environment Variables" on page 44
Create a service object of either type.	"Creating a Database Service Object" on page 58

Task	Reference
Use a DBResourceMgr service object.	“When to Use a DBResourceMgr Service Object” on page 57 “Connecting with a DBResourceMgr Service Object” on page 66
Use a DBSession service object.	“When to Use a DBSession Service Object” on page 57 “Connecting with a DBSession Service Object” on page 65
Dynamically select a database to connect to.	“Dynamically Choosing a Database Vendor” on page 69
Run iPlanet UDS using a standalone database.	“Using a Local Database when Running iPlanet UDS Standalone” on page 52
Write TOOL SQL statements to query and update database data.	“Using TOOL Statements to Query Data” on page 98
Write methods on DBSession class to generate SQL statements at application runtime.	“Using iPlanet UDS Classes to Execute SQL” on page 118
Reconnect a dropped database session.	“Reconnecting to a Database Session” on page 75
Disconnect the session.	“Disconnecting a Database Session” on page 76

Database Vendor Notes

Following are some general notes on how the iPlanet UDS database interface works with specific databases. This type of database vendor-specific information is found throughout the manual; for example, see [“Environment Variables” on page 44](#).

DB2 Notes

In IBM’s DB2 Client/Server architecture, a DB2/6000 application executing through the DB2 CLI Driver can connect to the following:

- a local DB2/6000 database using its DB2 server
- a remote DB2/6000 database using a local DB2 Client Application Enabler and a remote DB2 Client Support, using DB2 networking services

When connecting from an iPlanet UDS application to a DB2/6000 database, the first DB2 configuration above is used. iPlanet UDS partitions the application to execute the DB2 service object (either a DBResourceMgr or DBSession) on the RS/6000 node where the DB2 database resides. iPlanet UDS application services are used to communicate between the iPlanet UDS/DB2 server partition and other application components. The actual DB2 connection, then, is a local connection between the iPlanet UDS server partition and the DB2/6000 server.

iPlanet UDS provides array interface support for DB2 and ODBC. This standard can improve network performance when transferring multiple rows between a database and a client. Using the array interface requires only one round trip between the database and client to transfer multiple rows, rather than one round trip per row.

Informix Notes

An Informix database is identified by both an Informix server name and a database name. If a server name is not specified, then the value of the INFORMIXSERVER environment variable is used. Note that, unlike native Informix tools such as *dbaccess*, when iPlanet UDS connects to Informix a database must be specified; server-only connections are not supported.

NOTE An iPlanet UDS server partition is limited to a single shared memory connection to a local Informix database.

Oracle Notes

An iPlanet UDS application that integrates with an Oracle RDBMS offers two communication models; both models require SQL*Net to be installed on any machine which will initiate a database connection (specifically, any machine that will run a DBSession object—for example, a DBSession service object).

The first model installs the iPlanet UDS server partition on the same node as the Oracle RDBMS; in this scenario, SQL*Net is not required on either node. This model provides efficient and fast communication between the application and the Oracle database.

The second model allows the Oracle RDBMS to reside on any machine in the network, provided that SQL*Net is running on that machine. This model provides more flexibility.

The node on which the database is installed can even be a different operating system and hardware platform than the node running the iPlanet UDS application, as long as both the iPlanet UDS server and the Oracle server support the same network transport protocol (TCP/IP, for example) and SQL*Net is installed on the iPlanet UDS server machine.

Rdb Notes

iPlanet UDS can access both local and remote Rdb databases from OpenVMS systems. iPlanet UDS can also access non-Rdb databases through Oracle's Database Integrator (DBI) family of products. Note that iPlanet UDS requires only the Rdb Run-Time license to access Rdb databases.

Because of the design of the iPlanet UDS architecture, the connectivity model departs from conventional Rdb connections. In the conventional configuration of a local, native Rdb application, the Rdb runtime system runs within the server partition and accesses Rdb data on the server node.

In a typical Rdb client-server configuration, SQL/Services enables the client application to access Rdb data on the server node; SQL/Services software is required on both sides of the connection. DECnet, TCP/IP, and AppleTalk communication protocols are supported between SQL/Services clients and servers.

NOTE In an iPlanet UDS client-server application accessing local Rdb data, communication between the client and server is handled by iPlanet UDS rather than SQL/Services. If the iPlanet UDS server partition is installed on the same node as the Rdb database, iPlanet UDS will make an efficient local database connection.

The Rdb database can be installed on an OpenVMS machine that does not have iPlanet UDS software installed. The iPlanet UDS database service object accesses the remote Rdb database using DECnet and an RDB\$REMOTE account. While less efficient than the configuration described above, this illustrates the flexibility available with the iPlanet UDS/Rdb interface.

Unsupported Database Features

Most database management systems extend the ANSI standard SQL language. Whenever possible, iPlanet UDS supports vendor-specific SQL extensions; however some extensions are inconsistent with the implementation of TOOL SQL and cannot be supported. These are indicated below.

- scroll cursors (cursors that can be fetched in a non-next fashion)

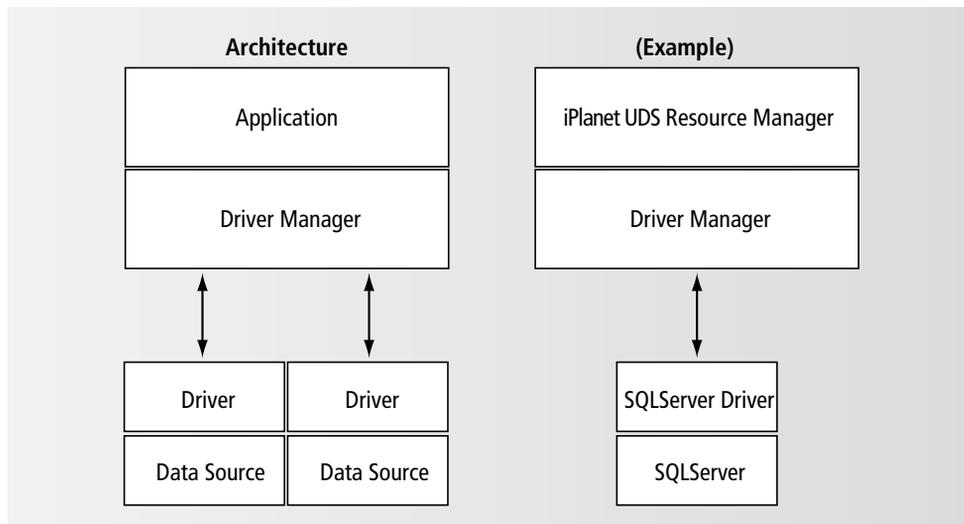
For Informix only, scroll cursors are supported through the DBSession class interface; scroll cursors are not supported for other databases.

- Informix *hold cursors* are not supported (cursors that remain open across transactions).

Using ODBC

Microsoft's Open Database Connectivity (ODBC) is a standard application programming interface (API) for accessing data stored in database management systems. Using ODBC, you can develop, compile, and ship an application that is database independent. Using standard SQL to access data, an ODBC application can access various databases without modification or recompilation. At runtime, dynamic-link libraries, called *database drivers*, are linked on demand to access a specific *data source* through a specific communication method.

The ODBC architecture has four main components, shown in [Figure 1-3](#) and described below.

Figure 1-3 ODBC Architecture

application Performs processing and calls ODBC functions to submit SQL statements that manipulate and/or retrieve data in the target database. When using iPlanet UDS with ODBC, the iPlanet UDS ODBC resource manager is the ODBC application. This component is the only one provided by iPlanet UDS; all other components must be in place for an iPlanet UDS application to use ODBC.

driver manager Loads database drivers on behalf of an application. This component allows ODBC applications to access different databases without recompilation. The driver manager is provided by Microsoft on Windows platforms.

database driver Processes ODBC function calls, submits SQL requests to a specific data source, and returns results to an application. Database drivers are provided by database vendors or other third parties.

data source The user's data, the associated operating system, database, and network platform (if any) used to access the DBMS.

Supported Data Sources

To see the official list of which data sources (and versions) iPlanet UDS certifies against which platforms, refer to the platform support matrix, which is available at <http://www.forte.com/support/platforms.html>.

Supported Database Drivers

In addition to a data source, you must have a supported database driver for that data source. The platform support matrix, at <http://www.forte.com/support/platforms.html>, lists certified database drivers for each platform.

Because various databases and database drivers offer a range of functionality, the ODBC interface defines conformance levels to categorize the level of support each driver and database provides. iPlanet UDS applications require drivers to support:

- Level 1 API and Core SQL grammar, as defined in the ODBC 2.1 SDK.

Also, iPlanet UDS can use ODBC array interface support if the ODBC driver supports:

- SQL Extended Fetch, so that iPlanet UDS can perform an array fetch
- SQL Param Options, so that iPlanet UDS can perform array inserts

If the driver does not support these features, iPlanet UDS simply does not use the array fetch and insert features.

Using ODBC When Running iPlanet UDS in Standalone Model

The primary benefit of ODBC is that it allows you to access a database system when running iPlanet UDS standalone, rather than distributed. This can be useful in a number of situations, for both application developers and end users.

For example, an application developer can prototype an entire distributed iPlanet UDS application running standalone on Windows running against a SQL Server data source. Later, the developer could deploy the application in a different environment against other, possibly multiple, databases.

Another example is an end user who daily accesses a local database while running standalone, but periodically connects to a distributed environment to access a different database. Such a user could use ODBC when running standalone, regardless of the type of database he accesses when running distributed.

When the iPlanet UDS development/runtime system is run standalone, iPlanet UDS always assumes a Database Resource manager type of ODBC. For more information on using ODBC in standalone mode, refer to [“Using a Local Database when Running iPlanet UDS Standalone” on page 52](#).

Cautions on Using ODBC with iPlanet UDS

When you use ODBC with iPlanet UDS you should be aware that not all data sources or all drivers provide the same functionality. Depending upon which data source and driver you use, some database features may not be available.

On NT and Windows 95, your ODBC driver must be thread safe. iPlanet UDS exploits the native threading capability on these platforms and requires the ODBC driver to operate in a multithreaded environment.

If an iPlanet UDS application attempts to use a database feature that is not supported by a particular data source, iPlanet UDS generates an exception (a `DBUsageException` or a `DBResourceException`, depending upon the circumstances). Examples of features that may generate this type of exception include outer joins and database procedures.

Defining a Resource Manager

Before an iPlanet UDS application can access data in a relational database, you must create a *resource name* for the database. This chapter contains the following topics:

- the purpose of the database resource manager
- choosing a node for the resource manager
- defining a database resource name
- running the DynamicSQL example to test a resource name
- accessing a local database when running iPlanet UDS standalone

If you are an iPlanet UDS system manager, follow the instructions in this chapter to set up and test a resource name.

If you are an iPlanet UDS application developer, ask your system manager for valid resource names. Then, to access a database from an iPlanet UDS application, you must establish database connections as described in [Chapter 3, “Making a Database Connection.”](#)

About Resource Managers

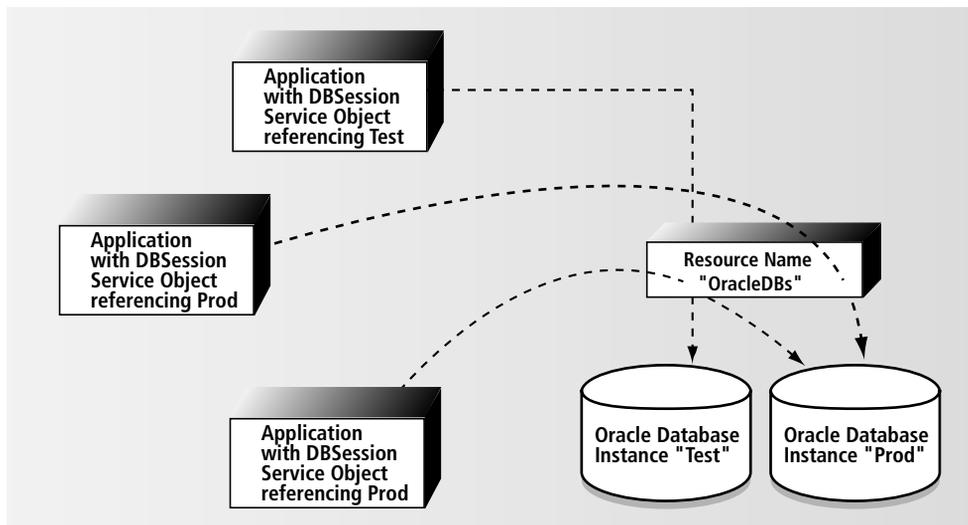
For a given database to be accessible by an iPlanet UDS application, a *resource manager* must be associated with the database. A resource manager simply indicates a *type* of database (as specified by the name of the database vendor, which indicates the database’s interface) that can be accessed from a particular *node*. Note that a resource manager does *not* uniquely identify a particular database instance, but only a database type, such as Oracle, Sybase, or ODBC.

The iPlanet UDS system manager creates a database *resource name* using the Environment Console. Each resource name is associated with a resource manager. See [“Defining a Resource Name” on page 42](#) to see an example of the dialog window. iPlanet UDS application developers use the resource name in applications that need to access a database of that type (vendor). During partitioning, iPlanet UDS places the database service objects on the same node as the resource managers they use. Subsequently, all database connections to a given database are established through the resource name for that database.

A resource name contains no connection information (specifically, unique database name, username, and password). All connection information is provided either through a DBSession service object, or when connecting to a database using the ConnectDB method on the DBResourceMgr class. See [Chapter 3, “Making a Database Connection”](#) for a discussion of database connection options.

One resource name can be used by multiple database service objects. Each service object would reference the same resource name, but could use different database names when providing the connection string to initiate a database session. [Figure 2-1](#) shows an example of three different applications using the same database resource name but connecting to different databases.

Figure 2-1 Resource Name used by Two Databases



Distributed Development and Access The iPlanet UDS system manager defines a resource name for access by multiple application developers (and end users) running iPlanet UDS in distributed mode. To do so, the manager uses the Node Properties dialog of the Environment Console, and defines the resource name on a node on which the DBMS client interface resides. The node need not have the database itself installed, as long as the node can access the database (for example, any database-specific environment variables are set). [“Defining a Resource Name” on page 42](#) describes this procedure.

Stand-alone Development and Access Any iPlanet UDS application developer or end user running iPlanet UDS in standalone mode uses an ODBC resource name by default; all database connections are assumed to use ODBC. The name of the data source must be specified, but the resource name and resource manager are assumed. See [“Using a Local Database when Running iPlanet UDS Standalone” on page 52](#) for more information.

Choosing a Node for the Resource Name

You should define a resource name on a node that has guaranteed and efficient access to the target database system or ODBC data source. When choosing the node, keep the following in mind:

- The iPlanet UDS database interface can be considered the client side of a specific database vendor’s client-server support. This is the case even though iPlanet UDS generates server partitions for the DBSession or DBResourceMgr service objects.
- By default, iPlanet UDS partitioning locates service objects on the same node as the resource manager they use.
- Some database vendors provide client-server connectivity using standard networking protocols or fast transfer using shared memory when the client is on the same node as the DBMS server. It may be advantageous, therefore, to define a resource name on the same node that runs the DBMS server process(es).
- For databases accessed through the ODBC interface, the resource name must be defined on a Windows node that contains the required ODBC driver manager and drivers.
- It can sometimes be advantageous to load balance client-server activity by having clients run on a node (or nodes) other than that which runs the server process(es).

Defining a Resource Name

If you are running iPlanet UDS standalone using an ODBC data source, this section is irrelevant because you do not use an Environment Console in standalone mode. However, if you are a developer who runs both stand-alone and connected (accessing both local and central databases), then your iPlanet UDS system manager would complete these steps to define a resource name for a server-resident database.

After you choose a node on which to define the resource name, you can use either the Environment Console or the Escript utility to define a resource name.

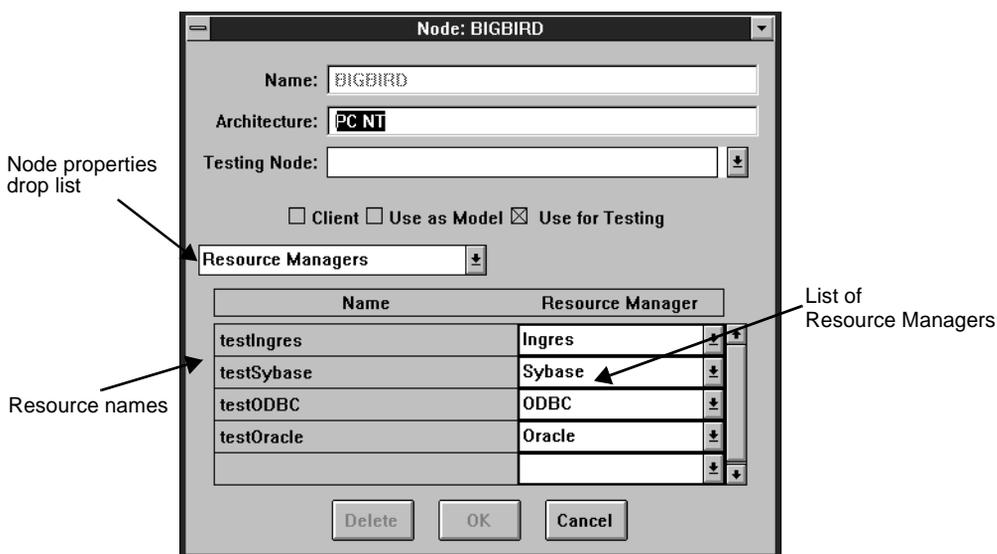
➤ **To define a resource name for a database using Environment Console**

1. Open the Environment Console.
2. Choose either the View > Node Topology or View > Node Outline command.
3. Lock the environment by clicking the Locked toggle.
4. Select the node on which you will define the resource name.

It is from this node that connections will be made to the database or ODBC data source. Therefore, this node must be running either the database system itself or the database client software.

- For this node, choose the Component > Properties... command.

The Node Properties dialog appears, as shown below. Make sure that Resource Managers, the default choice, appears in the node properties drop list.



- Enter any unique name to be the resource name for the database.

All resource names defined for a given environment must be unique within that environment (the name space is an iPlanet UDS environment).

You will use this name later, in the “Database Manager” field of the Database tab page of the Service Object Properties dialog (see [“Entering Database Information” on page 60](#)). Because the resource name is not physically mapped to any specific database, you can easily move the resource name to a different node, if desirable, without changing any iPlanet UDS code.

- Choose the appropriate database type (vendor) from the Resource Manager drop list.

The drop list varies by platform; it shows only database interfaces that iPlanet UDS supports for that node’s architecture. It does not reflect databases that are actually installed on the node.

- Save the changes to the environment using the File > Save Environment command.

9. Unlock the environment by clicking the Locked toggle. Exit the Environment Console.
10. If necessary, set the appropriate environment variables for the database, as described later in this chapter. On some platforms, notably VMS, other environmental considerations such as process privileges and quotas may apply.

Using Escript to Define a Resource Name

Rather than using the Environment Console, you can use the Escript command `AddExternalRM` to define a resource name. For more information on the Node Properties dialog, or the `AddExternalRM` command, refer to the *iPlanet UDS System Management Guide*.

Environment Variables

When you install or reinstall iPlanet UDS on a server node, you provide database-specific information that iPlanet UDS uses to automatically update all required environment variables for each database. So, you may not need to manually set any environment variables, because the installation does so for you. However, if you want to add a new database (that is, a new database instance from the same vendor, or a database from a different vendor) then you must update the environment variables so that iPlanet UDS can interact with the new database. Refer to the appropriate sections below to set any environment variables that apply to the database you are adding.

Environment variables at three different levels may affect your iPlanet UDS database application. You can set environment variables (or logical names) required or used by (1) the operating system, (2) by iPlanet UDS, or (3) by the database.

O/S environment variables You should set the UNIX `LIBPATH` environment variable to include the library directory for each database you will connect to on a UNIX node. Check either of the files `fortedef.sh` (for K or Bourne shell) or `fortedef.csh` (for C shell) and add the value for `LIBPATH` if necessary. The `LIBPATH` variable must be set before you start the iPlanet UDS Node Manager process.

The actual name for LIBPATH varies slightly by platform:

Platform	Name of LIBPATH to use
Compaq Tru 64 Unix, Solaris, DG, PTX	LD_LIBRARY_PATH
RS6000	LIBPATH
HP	SHLIB_PATH

iPlanet UDS environment variables iPlanet UDS requires no additional iPlanet UDS-specific environment variables to access databases. However, there are a few optional environment variables you can use to tune application performance (FORTE_DB_MAX_STATEMENTS and FORTE_DB_VENDORFLG are two that relate to databases; for more information on these variables see the iPlanet UDS online Help).

Vendor-specific environment variables All database vendor-specific environment variables must be set in the iPlanet UDS server process containing the database resource manager or session. Since iPlanet UDS server partitions inherit environment variables from their iPlanet UDS Node Manager, it is sufficient to set the database environment variables in the iPlanet UDS Node Manager process before any iPlanet UDS process is started.

The following table lists the database vendor-specific environment variables that must be set for a given vendor-type of database. These variables are described in following sections.

Database	Environment Variables
DB2	DB2INSTANCE and other variables specified in the files db2profile and db2cschrc
Informix	INFORMIXDIR, INFORMIXSERVER
Oracle	ORACLE_HOME, ORACLE_SID
Rdb	RDMS\$RUJ, SORTWORK, and RDMS\$DEBUG_FLAGS
Sybase	SYBASE, if root directory (interfaces file) is in non-standard location.

DB2 Variables

An iPlanet UDS Node Manager may access only one DB2/6000 instance. Each DB2/6000 instance has one of two example files, `db2profile` (for K shell) and `db2cschrc` (for C shell), which defines the environment variables necessary to access the instance (such as `DB2INSTANCE`). You should incorporate (add) the appropriate file into either your `.profile` file (for K shell) or `.cshrc` file (for C shell). See IBM's *DATABASE 2 AIX/6000 Installation Guide* for a description of these environment variables.

Informix Variables

Informix uses a number of environment variables to modify a given installation. Two variables are particularly important for iPlanet UDS applications:

INFORMIXDIR A required variable that specifies the top directory of the node's Informix product installation.

INFORMIXSERVER A required variable that can be used to specify a default Informix database server for database connections. When a connection `resourceName` parameter does not contain a server name component, the value set for `INFORMIXSERVER` is used.

Environment Variable	Platforms	Setting
INFORMIXSERVER and INFORMIXDIR	Solaris, Alpha OSF, DG Intel, PTX Intel	LD_LIBRARY_PATH should include \$INFORMIXDIR/lib:\$INFORMIXDIR/lib/esql
	RS6000	LIBPATH should include \$INFORMIXDIR/lib:\$INFORMIXDIR/lib/esql
	HP9000	SHLIB_PATH should include \$INFORMIXDIR/lib:\$INFORMIXDIR/lib/esql

Oracle Variables

ORACLE_HOME During the Oracle installation, you set the environment variable `ORACLE_HOME` to indicate the home directory for the Oracle installation. `SQL*Net V2` uses this environment variable to locate the file `tnsnames.ora`.

ORACLE_SID You should also set the environment variable `ORACLE_SID` to the instance of Oracle to which you are going to connect. Local connections to an Oracle database require `ORACLE_SID` to be set.

Rdb Variables

While there are no VMS logical names required to access Rdb databases from iPlanet UDS, some Rdb-specific logical names may be useful. Logical names can be used for two distinct purposes:

- Rdb configuration and load balancing, using the `RDMS$RUJ`, `SORTWORK`, and `RDMS$DEBUG_FLAGS` logical names.
- Naming databases, by defining a logical name that translates to an Rdb database specification. This logical name can then be specified in the `ConnectDB` method of the `DBResourceMgr` class (see [“Connecting to a Database” on page 55](#)). Using a logical name for a database makes your iPlanet UDS code and service object definitions more portable.

For an iPlanet UDS/Rdb session to process a logical name, the name must be “visible” to the Rdb server process created by the iPlanet UDS Node Manager. The recommended procedure is to define the name in the iPlanet UDS “global logical name table” by adding the definition of the logical name to the `FORTE_STARTUP` procedure. See the section on installing on Alpha OpenVMS in the *iPlanet UDS System Installation Guide* for more details.

Sybase Variables

During the Sybase installation, a file named “interfaces” is created in the root directory for Sybase. By default, this directory is the login directory for the Sybase account. Sybase CB-Library must be able to find the interfaces file to translate the server name to an appropriate network address. If the Sybase root directory is not in the login directory for the Sybase account, then you must use the environment variable `SYBASE` (a logical name on VMS) to tell Sybase where to find the interfaces file.

Removing a Resource Name

You might wish to remove a resource name from one node in an environment so that you can redefine it on another node. You might do this if you moved your database to a more powerful machine. You could redefine the resource name to use a resource manager on the newer machine; this would cause iPlanet UDS to place the servers that use that resource manager on the newer machine during partitioning.

To remove a resource name, complete steps 1 through 4 as described in [“Defining a Resource Name” on page 42](#). Then, highlight the resource name you wish to remove and click the Delete button.

Testing a Resource Name with the DynamicSQL Example

You can use the iPlanet UDS sample application called DynamicSQL to test a new or existing resource name. This section contains instructions for running DynamicSQL. When the instructions differ by database vendor, follow the instructions that pertain to your particular vendor.

The DynamicSQL example does not create any new tables or add any data to your database. It simply opens an iPlanet UDS window through which you can enter any SQL statement using tables and data that already exist in the database.

► To test a resource name

1. Using the vendor’s utilities, verify that the database (or ODBC data source) is accessible. After you complete this step, your database should also be accessible from iPlanet UDS.

Database Vendor	To verify that the database is accessible:
DB2	Login to the node. Define the environment variables appropriate for accessing the local DB2/6000 database “dbname.” Run the DB2 command line processor (“db2”) and use the following command: <code>connect to dbname user uname using upass</code>
Informix	Login to the node using the uname account. Define the INFORMIXDIR environment variable, and if appropriate, INFORMIXSERVER. Use the dbaccess command to connect to the Informix database.

Database Vendor	To verify that the database is accessible:
ODBC	Login to the node and use the database-specific interface to test database access. For example, if you are using a Microsoft SQL Server data source, you would run the isql command line processor and connect to the dbname database, using <i>uname</i> and <i>upass</i> .
Oracle	Use the following command: <pre>sqlplus scott/tiger@servicename</pre> When you see a SQL> prompt, type Quit to exit the program.
RDB	Use the following command: <pre>\$ RMU /SHOW VERSION</pre> You should see the message <pre>"Executing RMU for DEC Rdb Vx.y."</pre> If your database node has an Rdb Interactive or Development license, then you can verify that the Rdb database is accessible using the SQL\$ Interactive SQL program. For example, if the database filename associated with MYDBNAME is DKAO:[DATABASES]FINCHES, then enter: <pre>\$ RUN SYS\$SYSTEM:SQL\$ SQL> ATTACH 'FILENAME DKAO:[DATABASES]FINCHES' : SQL> SHOW TABLES SQL> EXIT</pre>
Sybase	Use the following command: <pre>isql -S DBname -U uname -P upass</pre>

2. Create a resource name for the database (vendor) of your choice.

If you already have an existing resource name, go to [Step 3](#).

If you do not have an existing resource name, follow the instructions in the previous section. The DynamicSQL example assumes that the resource name is "AnyDBMgr." You can either use AnyDBMgr, or modify the reference to the resource name to match your name (described in the next step).

3. Create the DynamicSQL project. Open the Repository Workshop and import the two files:

```
$FORTE_ROOT/install/examples/frame/utility.pex
```

```
$FORTE_ROOT/install/examples/database/dynsql.pex
```

If you have an existing resource name, you must edit the service object for this example to use your resource name. In the Project Workshop for the DynamicSQL project, double-click on the AnyDBMgr service object. In the Service Object Properties dialog (General tab page), change the name in the Database Manager field from "AnyDBMgr" to match your resource name.

4. Run the DynamicSQL application.

The application input and output appear in the console/trace window, so you should open the window if it is iconized.

5. You are prompted for a database name, a username and a password. Respond to the prompts using the format that corresponds to your database vendor, as described in the following table.

Database Vendor	Database	Username	Password
DB2	<i>dbname</i> The name of the local DB2 database, as specified on a DB2/6000 SQL CONNECT statement. This database must be accessible from the local instance defined by the DB2INSTANCE environment variable. Example: TestDB2DB	<i>uname</i> An authorization name associated with a DB2/6000 SQL CONNECT statement. <i>uname</i> must be an authorized login user name on the AIX system.	<i>upass</i> The valid AIX password for <i>uname</i> .
Informix	<i>dbname@dbserver</i> The name of the Informix server and database. or <i>dbname</i> The name of the Informix database. The translation of the INFORMIXSERVER environment variable is used for the Informix server. Example: testinfdb or testinfdb@HQ3	<i>uname</i> A valid UNIX account on the iPlanet UDS server node.	<i>upass</i> The password associated with the <i>uname</i> account.

Database Vendor	Database	Username	Password
ODBC	<p><i>dbname</i></p> <p>The name of a valid ODBC data source.</p> <p>Example: MyLocalDB</p>	<p><i>uname</i></p> <p><i>uname</i> must follow the conventions related to the database to which you are connecting.</p>	<p><i>upass</i></p> <p><i>upass</i> must follow the conventions related to the database to which you are connecting.</p>
Oracle	<p>@[<i>TNSResourceName</i>]</p> <p>A database name, as specified in the SQL*Net V2 file ORACLE_HOME/network/admin/tnsnames.ora</p> <p>If <i>TNSResourceName</i> is omitted, then the connection is to a local database and the environment variable ORACLE_SID must be set. <i>TNSResourceName</i> is case-sensitive.</p> <p>Example: @testora or @</p>	<p><i>uname</i></p> <p>A valid username for the current Oracle database instance.</p>	<p><i>upass</i></p> <p>The password for the Oracle username.</p>
Rdb	<p>RDB.<i>dbname</i></p> <p><i>dbname</i> is the Rdb database filename specification used to attach to the database from the current node. A logical name can be used only if it is “visible” to the iPlanet UDS process associated with the RDBMGR service.</p> <p>The Rdb connection established by iPlanet UDS runs under the persona of the specified VMS user. This persona includes the username, UIC, authorized privileges and (new to iPlanet UDS Release 3), any general identifiers associated with the user. The Rdb database should have appropriate access granted to the user's UIC and/or identifiers.</p> <p>Example: RDB.MYDBNAME</p>	<p><i>uname</i></p> <p>A valid VMS username.</p>	<p><i>upass</i></p> <p>The password associated with the <i>uname</i> account.</p>

Database Vendor	Database	Username	Password
Sybase	resource. <i>dbname</i> resource is the name of the resource in the Sybase interfaces file. dbname is the name of the database for that server. Example: SYBASE.master	<i>uname</i> A valid Sybase username.	<i>upass</i> The password associated with the <i>uname</i> account.

- To use the application, enter SQL select statements for tables that you know to exist in the database. Type Quit to exit.

(Appendix A, “Database Example Applications” also contains instructions for installing and running the DynamicSQL example.)

Using a Local Database when Running iPlanet UDS Standalone

You can run iPlanet UDS in standalone mode on Windows platforms iPlanet UDS (as opposed to distributed mode) and connect to a local database using the ODBC database interface. Using this configuration you can run standalone iPlanet UDS database applications, or develop portable database applications that can later be deployed in distributed systems running on multiple platforms and accessing multiple databases from different vendors.

When you run iPlanet UDS standalone, an ODBC resource manager is always assumed. iPlanet UDS automatically assumes a resource manager type of ODBC for any DBSessions or DBResourceMgr service objects; you simply need to enter valid, local data source name when you define a database service object. For additional information about running standalone refer to “[Entering Database Information](#)” on page 60.

Setting up an iPlanet UDS ODBC Data Source on Windows

This section contains instructions for making an ODBC data source available to an iPlanet UDS system running standalone on any Windows platform.

NOTE These instructions assume that you have already installed the underlying data source (database). They allow you to define a data source that is running either locally on the same machine or remotely; in either case iPlanet UDS standalone can use it as a data source.

► **To define a new ODBC data source**

1. Install the ODBC Driver Manager (this is a component of ODBC, not the driver itself).
2. Install the appropriate database driver for each target database (each database that you wish ODBC to access for your iPlanet UDS application).

For a review of the ODBC architecture, showing the Driver Manager and database drivers, see [Figure 1-3 on page 36](#). Also, refer to the *ODBC 2.0 Programmer's Reference and SDK Guide* for additional information about installing these components.

3. Open the ODBC Administrator from the ODBC Control Panel. Click on the Add button. You will see a list of installed drivers, from which you select one for the new data source.
4. Choose the appropriate driver for your target data source and click OK.

To see a list of installed drivers, you simply click on the Drivers button on the ODBC Administrator dialog.

For example, to create a data source for a local SQL Server database, choose SQL Server as the driver type.

5. In the next dialog that appears, the ODBC Data Source Setup, enter information about the Data Source.

You can specify any database accessible to this node, whether local or remote. You can even specify a flat file as a data source using ODBC.

Enter any arbitrary name; you will use this name later as the Database Name in the Service Object Properties sheets, when specifying connection information.

Enter appropriate information for the Server, the Network Address and Network Library. At this point you can specify either a local or remote data source.

When you click OK, you have added a new data source.

After an ODBC data source has been defined to ODBC, there are a few additional steps required to make the data source available to iPlanet UDS users and developers.

➤ **To make the data source available to iPlanet UDS users**

1. The data source (for example, a database) must be running. From the program group for the database, start the database service by double clicking on the database icon and selecting Start.
2. (If you want iPlanet UDS users who are running distributed to be able to access the ODBC data source) In the Environment Console, add an ODBC resource manager for the node on which you defined the data source.

(For detailed instructions see [“Defining a Resource Name” on page 42.](#))

Using an Existing Data Source

To see the names of existing ODBC data sources, open the ODBC Administrator. You will see a display containing the names of valid data sources with the drivers they use.

To see the available drivers, press the Drivers button in the ODBC Administrator.

Making a Database Connection

This chapter describes options you have for connecting to a database. Many applications will use a service object to initiate a database connection. Both types of database service objects, `DBSession` and `DBResourceMgr`, allow you to connect to a database, establish a database session, and query and update database data.

You can also connect to a database without using a database service object. This process is also described in this chapter.

This chapter contains the following topics:

- choosing between the two types of database service objects
- creating a service object and connecting to a database
- connecting to a database without a service object
- using connection options, such as creating a default `DBSession` service object (a default user name and password) and using variable user names and passwords
- resuming dropped database connections

Connecting to a Database

You connect to the database by establishing a database session. A database session is a connection to a particular database (using an iPlanet UDS database resource name) through which you execute an isolated set of statements.

You can open and use any number of sessions to a single database, subject to operating system and database limitations. For example, to avoid creating a long-term transaction when an end user is browsing through the database and making changes, you can use two different database sessions. The first session can

open a cursor for read only, and the second session can use short transactions to make any database updates that the end user requests. You can also use sessions to access more than one database or type of database in the same iPlanet UDS application.

To initiate a database session you specify connection information (a database and a valid username and password). There are three ways to start a database session:

- using a DBSession service object (based upon the DBSession class)
- using a DBResourceMgr service object (based upon the DBResourceMgr class)
- using no service object

The following sections describe how to start a database session.

Every database connection (session) is represented by an object of class DBSession. When you invoke the ConnectDB method on a DBResourceMgr service object, the ConnectDB method returns a DBSession object. In fact, iPlanet UDS uses the same ConnectDB method for every database connection, but the method is transparent when used with a DBSession service object or without any service object at all. Subsequently, your application can execute TOOL SQL statements or invoke methods from the GenericDBMS library during a database session; both cases cause iPlanet UDS to invoke methods on the DBSession class on the current DBSession object.

NOTE Every SQL statement issued against a database is associated with an on session clause that identifies an active database session. You can use a service object name or reference to an object of the DBSession class. You can also omit the on session clause to specify the default session called DefaultDBSession (see [“Creating a DefaultDBSession Service Object” on page 73](#)).

Using Database Service Objects

Typically you will use either a DBSession service object or a DBResourceMgr service object to initiate connections to a database, although you can make connections without a service object (see [“When You Should Not Use a Service Object” on page 58](#)). The two service objects are used in different circumstances, as described below.

(See [“Creating a Database Service Object” on page 58](#) for instructions for defining either type of service object.)

When to Use a DBSession Service Object

DBSession service objects are perhaps the most convenient way to access a database, but offer somewhat less flexibility than the DBResourceMgr service. You can use a DBSession service object if the following are true:

- You can pre-specify the connection information (resource name, database name, user name and password) when you define the service object.

(You can override user information using environment variables; see [“Using Variable User Names and Passwords”](#) on page 71.)
- Your application does not require a dynamic number of database sessions. While there can be multiple sessions, the number of sessions cannot change at runtime.

The DBSession service object has the following advantage:

- It automatically connects to and disconnects from the specified database, using the pre-specified connection information, at application start-up and shutdown.

The DBSession service object has the following disadvantages:

- You must provide the user name/password and server/database name when you partition the application that contains the service object. (You may not specify these items after partitioning or when the application is running.)
- If the service object’s visibility is set to environment, then many clients share the same DBSession service object; thus, you cannot use individual users’ accounts for the sessions. Instead, you must use a generic account that stays the same across transaction boundaries.

When to Use a DBResourceMgr Service Object

You must use a DBResourceMgr service object if you want your application to obtain connection information at runtime.

The DBResourceMgr service object has the following advantages:

- You can invoke the ConnectDB method multiple times to start multiple database sessions.
- You can use different connection information for each database session.

The primary disadvantage to using a DBResourceMgr service object is that you must explicitly manage the connects and disconnects in your program. Some vendor-specific limitations do exist for DBResourceMgr service objects:

Rdb All concurrent sessions within a DBResourceMgr must use the same user name.

Oracle on VMS Only one session can be directly connected to the Oracle RDBMS; other concurrent sessions must be connected via the SQL*Net mailbox driver.

Informix Only a single shared memory connection may be active in an iPlanet UDS partition. Multiple network connections to Informix databases are fully supported.

When You Should Not Use a Service Object

In many applications, all database connections are made through a database service object. However, a characteristic of any database service object is that it is associated with one database when it is defined, and can only be used to connect to that particular database. This characteristic of service objects means that any application that must be able to connect to multiple databases cannot use either type of service object. For example, many applications might reasonably prompt users to specify which database they wish to connect to when they start the application.

You can write an application that is not bound to one database. Rather than using a service object to initiate a database session, you invoke the Connect method on the DBSession class, after you have determined which vendor-type of database the user will connect to. The vendor type of library must be determined first, because all objects and methods that iPlanet UDS uses during a given database session are actually vendor-specific.

See [“Connecting to a Database without a Service Object” on page 68](#) for detailed instructions on initiating database sessions without service objects.

Creating a Database Service Object

A project can contain any number of DBResourceMgr or DBSession service objects. To create a new service object, use the Project Workshop (see *A Guide to the iPlanet UDS Workshops* for details). Then use the New > Service Object command. The Service Object dialog appears, as shown in [Figure 3-1](#).

Figure 3-1 New Service Object Dialog

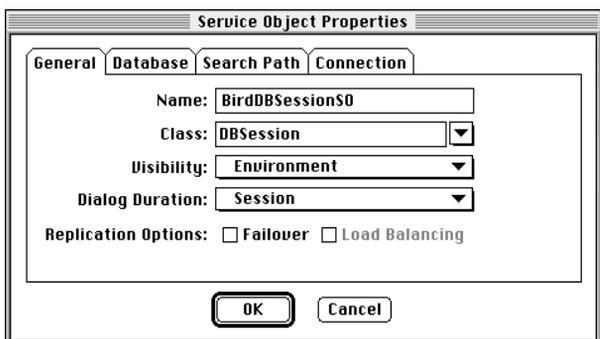
Specify a name for the service object and indicate whether it is a DBSession or DBResourceMgr service object. The Service Object Properties dialog appears next.

The Service Object Properties dialog uses tab pages to define the properties of one service object. A DBResourceManager uses three tab pages while a DBSession service object uses four pages.

To see more information on using the iPlanet UDS workshops, or the procedures and dialogs you use to create service objects, refer to *A Guide to the iPlanet UDS Workshops*.

Entering General Properties

Both types of service object use the General Properties tab page, shown in [Figure 3-2](#).

Figure 3-2 General Properties Tab Page

Visibility The default visibility of database service objects is environment-visible.

Dialog Duration If you want the service object to be replicated for failover, load balancing, or both, then you must use either Message or Transaction dialog duration; you cannot use Session duration with replication.

Replication Options Check one or both of the fields “Load Balancing” or “Failover” if you wish the service object to have these characteristics. For more information see the *iPlanet UDS Programming Guide*.

Entering Database Information

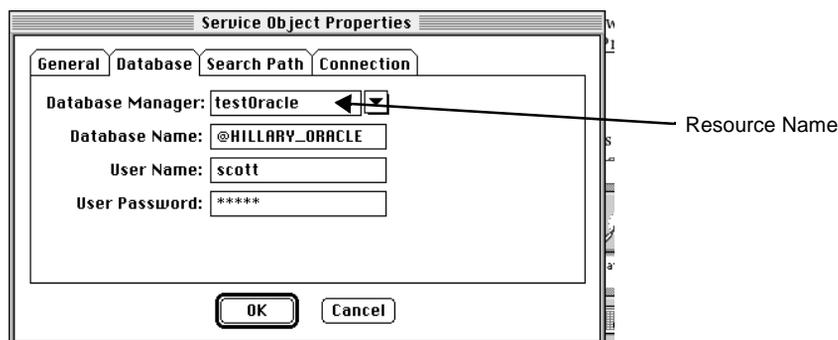
Both types of service object use the Database tab page, although the fields differ for the two types of objects.

For both types of service objects you specify the resource name to be used for database connections initiated by this service object. Use the drop list to enter this information in the “Database Manager” field.

For a DBResourceMgr service object, this tab page shows only the “Database Manager” field. Only the database manager information is required, because this type of service object allows different user information to be specified whenever a database connection is initiated.

For a DBSession service object, three additional fields appear as shown in **Figure 3-3**. You must enter user information for a DBSession service object, as it uses the same user information whenever it requests a database connection.

Figure 3-3 Database Tab Page



Database Manager The drop list for this field shows the resource names that have been defined by the iPlanet UDS system manager (as described in “[Defining a Resource Name](#)” on page 42.) Select the resource name for the type (vendor) of database with which this service object will interact.

If you are running iPlanet UDS stand-alone, the drop list is empty. It does not matter what you enter, because iPlanet UDS assumes an ODBC resource manager; however, you must enter something. This name is not case sensitive; in fact, you do not use it anywhere else.

Database Name Enter the name of the actual database with which the service object will interact.

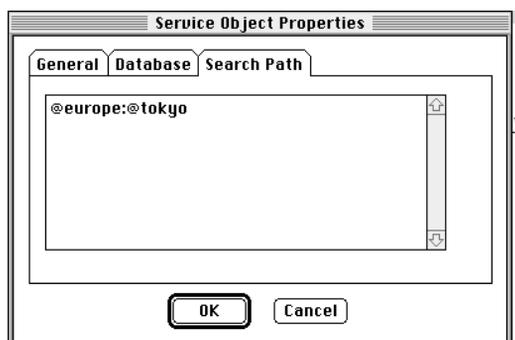
Vendor	Format	Description
DB2	dbname	<i>dbname</i> is the name of the DB2 database as known to the DB2 instance (defined by the DB2INSTANCE environment variable) running on AIX.
Informix	<i>dbname</i> [@ <i>dbservername</i>]	If <i>dbservername</i> is not specified, then the Informix environment variable INFORMIXSERVER must be set to identify the default server.
Oracle	@[TNSResourceName]	The <i>TNSResourceName</i> , including the @ symbol, is passed, unedited, to the Oracle connection routine. If only the "@" symbol is specified, then a local connection is made to the instance associated with the Environment Manager's ORACLE_SID value.
ODBC	data source name	Specifies the data the application will access, and its associated operating system, DBMS, and network platform (if any). Data source names are case sensitive.
Rdb	<i>DBtype.databasename</i>	<i>DBtype</i> is the type of database to which you are connecting (see the Rdb table in the iPlanet UDS online Help) <i>databasename</i> is the location of the database.
Sybase	<i>servername</i> [. <i>databasename</i>]	<i>servername</i> is the SYBASE domain server name. If you do not specify the database name, iPlanet UDS uses the Sybase default database for the user.

Username and Password Enter a valid username and password for the target database. Case sensitivity of these two fields depends upon the target database.

Entering Search Path Information

Both types of service object use the Search Path tab page, as shown in [Figure 3-4](#). Refer to *A Guide to the iPlanet UDS Workshops* for a discussion of how the search path is specified and used.

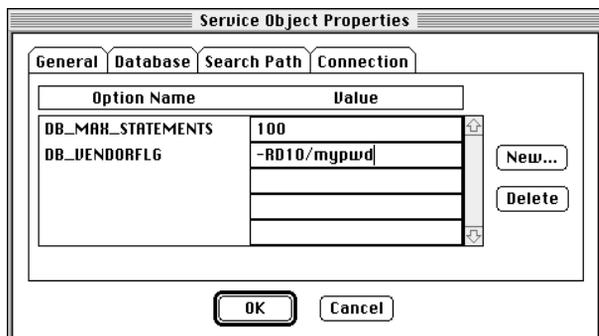
Figure 3-4 Search Path Tab Page



Entering Connection Information

Only a DBSession service object uses the Connection tab page, shown in [Figure 3-5](#). This page is entirely optional. Use this page if you want to specify connection options for any database connection to be made by the DBSession service object.

To specify each connection option you use a name-value pair. To see a list of valid options, refer to the description of the optionList parameter for the Connect method (see the iPlanet UDS online Help).

Figure 3-5 Connection Tab Page

You can specify Ingres connection flags by using the DB_VENDORFLG option. For example, to specify a role to use for an Ingres connection, you would enter DB_VENDORFLG under OptionName, and *-rroleID/password* under Value.

Specifying Connection Options in a Method

While only the DBSession service object allows you to enter connection options using the Connection tab page, you can enter connection options if you connect using either the ConnectDB method on DBResourceMgr or the Connect method on DBSession. An example follows:

```
options : Array of NamedElement = new;
opt : NamedElement = new (Name= 'DB_VENDORFLG',
    Object=TextData(Value=' -RroleID/password' ));
options.Append(opt);
myResource.ConnectDB
    (resourceName= 'IngresDB',
    userName='ingres', userPassword="",
    optionList = options);
```

Optimizing Service Object Performance

This section briefly touches upon some considerations that impact performance of your application. It does not discuss a number of designs that can be used to optimize performance of typical types of database applications.

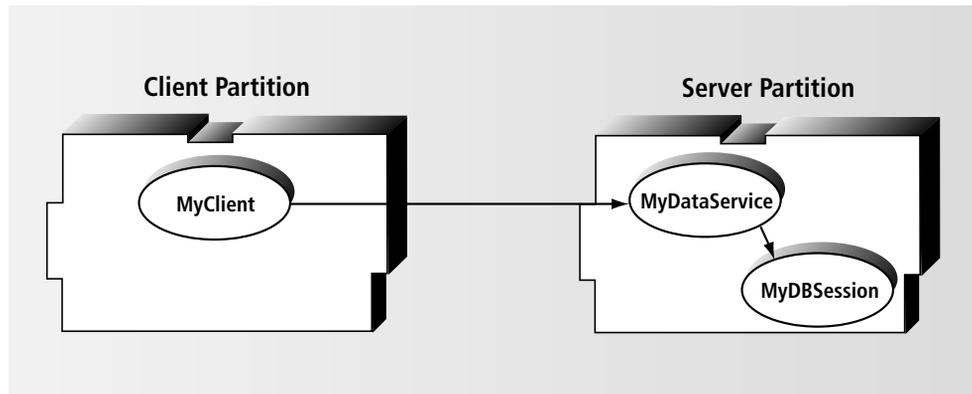
The location of a DBResourceMgr service object or DBSession service object does impact the performance of your application. By default, iPlanet UDS automatically places these service objects in a server partition on a node that supports the database vendor being used. However, you can easily change iPlanet UDS default partitioning.

The most critical factor in determining which partition to use for a service object is “Where is the *user* of the DBSession object with respect to the session?” For best performance the user should be placed in the same partition as the DBSession object.

While the DBSession class is distributed and can be referenced remotely, doing so will degrade performance. There may be circumstances when an application designer must use a remote DBSession object, but the overall application design must take this into account. Remote access is preferably used primarily during development and testing phases.

This is not to say that data cannot originate from, or be destined for, a remote partition. In fact, this is a typical scenario for a multi-tiered application. The best way to do this is to place a service object responsible for the data manipulation in the same partition as the database and channel access to the database through that service object. **Figure 3-6** illustrates this:

Figure 3-6 Partitioning with Intermediate Data Service on Same Partition as Session Service



Any access to data in the database managed by MyDBSession from MyClient must go through MyDataService. Access to DBSession is from MyDataService which is *local* to MyDBSession. This design optimizes performance, since only the data is sent between the client and server partitions in a single client/server interaction. Whereas, if MyClient accesses MyDBSession directly, the DBSession conversation takes place remotely and requires many more client/server interactions. This design has other benefits:

- MyDataService can do any processing of the data appropriate at this level, potentially resulting in less data to move between the client and server.
- The location of MyDataService can be moved at partitioning time if the database is moved.
- MyDataService can be replicated for fault tolerance or load balancing.

Visibility of the Service Object

If you use an architecture like that shown in [Figure 3-6](#) you should also set the visibility for the MyDBSession service object to “user” visibility. User visibility means that the service object is only known within its containing partition, not outside. This prevents unintended distributed access; MyClient cannot access or use the MyDBSession service.

On the other hand, the MyDataService object, as the public interface to persistent data, requires “environment” visibility so that it can be seen by MyClient (and any other client). MyDataService would in turn implement this persistence using MyDBSession as the data store.

Making a Database Connection

The following sections describe the standard ways to connect to a database, using either type of database service objects, or connecting to a database specified dynamically at runtime.

Connecting with a DBSession Service Object

A DBSession service object implicitly connects to and disconnects from the database. After you have created a DBSession service object, you simply reference that service object name in a SQL on session clause each time you access the database.

➤ **To start a database session using a DBSession service object**

1. Define a DBSession service object. (See [“Creating a Database Service Object” on page 58.](#))
2. Use the service object name in the on session clause of your SQL statements to specify the database connection the statement should use.

NOTE In the on session clause you can use a service object name or reference to an object of the DBSession class. If you omit the on session clause, the default session called DefaultDBSession is used (see [“Creating a DefaultDBSession Service Object” on page 73.](#))

For example:

```
sql execute immediate
'create table ArtistTab(Name char(30) not null, Country char(30)
not null)'
on session MySession;
```

The session connects to the database using the user name, password, and database name that were specified when the DBSession service object was created in the Project Workshop. You can override this information at several points, using either the Service Object Properties dialog in the Partition Workshop or the Environment Console.

Connecting with a DBResourceMgr Service Object

One benefit of using a DBResourceMgr service object is that database connection information can be provided at runtime.

► **To start a database session using a DBResourceMgr service object**

1. Define a DBResourceMgr service object. (See “Creating a Database Service Object” on page 58.)
2. Explicitly invoke the ConnectDB method on the DBResourceMgr service object in your TOOL code.

Your program may prompt users for database name, user name, and password, as shown in Figure 3-7. If you do not prompt the user for userName, userPassword, and resourceName, you must enter the actual values to be used for the ConnectDB parameters.

The ConnectDB method returns a DBSession object, which you will reference when you access the database.

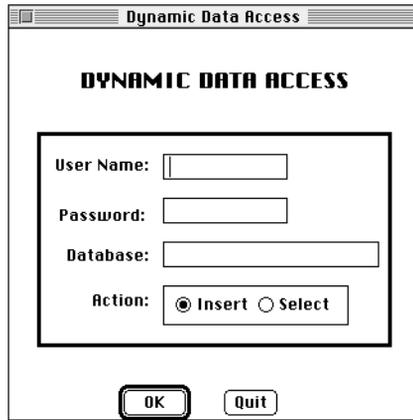
```
MySession : DBSession;
MySession = DBResMgr.ConnectDB(resourceName = DBName,
                               userName = UserName,
                               userPassword = Password);
```

3. Reference this DBSession object in the on session clause of your SQL statements.
4. If desired, invoke the Reconnect method (of the DBSession class) on a DBRemoteAccessException.
5. Explicitly disconnect each DBSession from the database.
6. Set the DBSession object to NIL for optimal memory management.

When the partition that issued the ConnectDB exits, any active DBSessions are automatically disconnected.

Getting connection information The iPlanet UDS DynamicDataAccess sample application demonstrates obtaining database and user information before starting a database session. It displays a login screen, shown in Figure 3-7, in which you enter a user name, password, and the database to which you want to connect. Once this information is entered, the application invokes the ConnectDB method on the DBResourceMgr service object and connects to the database using the login input as parameters.

Figure 3-7 Login screen for DynamicData Access Example



For more information about installing and running the Dynamic Data Access example, see [“DynamicDataAccess” on page 163](#).

Connecting to a Database without a Service Object

When you use a DBSession service object or a DBResourceMgr service object, the service object is associated with a specific database resource name during partitioning. For a DBSession service object the user name and password are also pre-specified. So, when you access a database using either type of service object, you must access a predetermined database. If you require more flexibility (specifically, you would like to choose which database to connect to) you do not need to use a service object to initiate a database connection.

The Connect method of the DBSession class allows you to specify, at runtime, connection information (database vendor, database, and user information) to be used to establish the database session. The Connect method also provides some additional security over the use of an environment visible service object.

By using the Connect method and dynamically loading libraries, you can dynamically start a database session without a service object.

The Connect method has the same parameters as, and functions similarly to, the ConnectDB method in the DBResourceMgr class. However, the ConnectDB method creates and returns a DBSession object, whereas the Connect method works on an existing DBSession object.

NOTE The Connect method does not function on a DBSession service object. The Connect method will only work on a DBSession object created at runtime, either directly as described below or via the ConnectDB method on the DBResourceMgr class.

For reference information on the Connect method, see the iPlanet UDS online Help.

Dynamically Choosing a Database Vendor

This section describes how to use the Connect method to create and connect a database session. This procedure allows you to indicate the database vendor to be used for the connection (in contrast to when you use a service object).

► To create and connect through a DBSession object

1. Load the library that corresponds to the desired database vendor using the FindLibrary method on the Partition object.

To invoke the FindLibrary method for a database vendor, you must know both the library name and a file key. The vendor's library must be available on the current node. Valid values are shown in the table below.

Database Vendor	LibName	FileKey	ClassName
DB2	DB2	D2	qqdb_DB2Session
Informix	Informix	IX	qqdb_InformixSession
ODBC	ODBC	OD	qqdb_OdbcSession
Oracle	Oracle	OR	qqdb_OracleSession
RDB	Rdb	RD	qqdb_RDBSession
Sybase	Sybase	SY	qqdb_SybaseSession

2. Create a DBSession object of the desired database vendor subclass.

For this step you must know the name of the DBSession subclass for the desired database vendor. These classes are shown in column 4 of the table above.

3. Connect the DBSession object to the database using the Connect method.

The TOOL code that follows specifies a particular database vendor and creates a vendor-specific DBSession object. In this code you would substitute actual values for the three variables *LibName*, *FileKey*, and *ClassName*; these parameters take string or TextData values.

```
DBVendorLibrary : Library;
DBVendorClassType : ClassType;
Session : DBSession;
DBVendorLibrary =
    task.part.FindLibrary(LibName, NIL, 0, FileKey, 4);
DBVendorClassType =
    DBVendorLibrary.FindClass(ClassName);
Session = DBSession(DBVendorClassType.InstanceAlloc());
-- then invoke Session.Connect ...
```

Initially the DBSession object is not connected to the database; the Connect method establishes the actual connection. For example, to establish a connection to an Oracle database you would use the following code:

```
DBVendorLibrary : Library;
DBVendorClassType : ClassType;
Session : DBSession;
DBVendorLibrary =
    task.part.FindLibrary('Oracle', NIL, 0, 'OR', 4);
DBVendorClassType =
    DBVendorLibrary.FindClass('qqdb_OracleSession');
Session = DBSession(DBVendorClassType.InstanceAlloc());
Session.Connect(resourceName = '${MY_DATABASE}',
                userName = 'scott',
                userPassword = 'tiger');
```

Just as with database service objects, you can use environment variables for the name of the database, as shown in the example (as well as for the user name and password).

Because the DB2 library is a subclass of ODBC, you must invoke the FindLibrary method twice, first for ODBC and then for DB2. So the example above would require one additional call to use DB2:

```
DBVendorLibrary =
    task.part.FindLibrary('ODBC', NIL, 0, 'OD', 4);
DBVendorLibrary =
    task.part.FindLibrary('DB2', NIL, 0, 'D2', 4);
```

For more information, see the iPlanet UDS online Help.

Other Connection Information

This section describes connection options, such as specifying a default database session or using default connection information (user name and password). It also describes how to reconnect a dropped session or disconnect a session.

Using Variable User Names and Passwords

Although the DBSession service object requires a user name, password, and database name, you can use environment variables to provide this information at start-up time. You can also reference environment variables at some points in your TOOL code. When iPlanet UDS encounters a variable, it looks to the environment for the actual value, and replaces the variable accordingly in the TOOL code.

NOTE The server portions of the application (which include the DBSession start-up code) are started by the Node Manager process running on the node where the resource name is defined. Therefore, the environment variables need to be defined in the command shell that starts the Node Manager process. Also, depending on the operating system, it may not be possible to change the values being used by the Node Manager once it has started. This is the case on UNIX, for example.

You can use environment variables when defining your DBSession service object as a way to dynamically change the database name, user name, and password without altering your application—you simply change the value of the environment variable. Specify the values for these parameters as strings or TextData values containing a dollar sign, followed by an environment variable name surrounded by braces.

For example, you could use the following values for database name, user name, and user password when you define your DBSession service object in the Service Object Properties dialog:

- `${SITE_DATABASE}`
- `${SITE_USERNAME}`
- `${SITE_PASSWORD}`

Figure 3-8 Using Environment Variables in Service Object Definition

The screenshot shows a dialog box titled "Service Object Properties". It contains several fields and options:

- Name:** jenn
- Class:** DBSession (dropdown menu)
- Visibility:** Environment (dropdown menu)
- Dialog Duration:** Session (dropdown menu)
- Replication Options:** Failover Load Balancing
- Database Manager:** DocOracle (dropdown menu)
- Database Name:** \${SITE_DATABASE}
- User Name:** \${SITE_USERNAME}
- User Password:** *****
- Shared DBSession
- Buttons: Search Path..., OK, Cancel

As iPlanet UDS processes the values it checks for a dollar sign followed by a name enclosed in braces. If braces are found, iPlanet UDS checks for an environment variable with the same name. If variables are found, their values are substituted for database name, user name, and user password. For example, your environment could include the following definitions:

- `SITE_DATABASE = @Test2Oracle`
- `SITE_USERNAME = scott`
- `SITE_PASSWORD = emu`

This creates a session for scott/emu on the database Test2Oracle. You can create a connection for different users and databases by simply changing the environment variables.

You can also use this technique with a DBResourceMgr service object: simply use the variables as the values for the parameters of the ConnectDB method, as shown below:

```
self.Session = AnyDBMgr.ConnectDB(
    username = '${SITE_USERNAME}',
    password = '${SITE_PASSWORD}',
    ResourceName = '${SITE_DATABASE}');
```

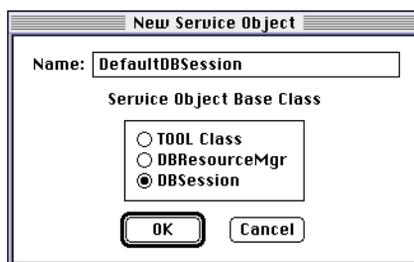
Creating a DefaultDBSession Service Object

You can establish a default database session for your SQL statements. A default database session allows you to issue a SQL statement without the on session clause. Using a default database session assures you that your program logic will always contain a reference to a database session. However, if you structure your code to use several sessions, you will need to change your TOOL code.

You can create a default database session for several scopes. You can create a default database session that will apply to an entire project, to a class and all of its associated methods, or to one method only. Each type is explained below.

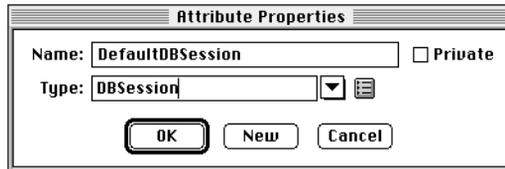
► To create a default database session for an entire project

1. Create a DBSession service object.
2. Name it DefaultDBSession, as shown below:



➤ **To create a default database session for a class and its associated methods**

1. Create a DBSession service object called AnotherDBSession.
2. Create a class called DataSelection.
3. Create an attribute for the above class called DefaultDBSession, as shown below.



4. In the Init method of the class, assign the DefaultDBSession attribute to the existing service object, as shown below.

```
super.init;  
DefaultDBSession = AnotherDBSession;
```

5. Create a method DoSelect for the DataSelection class that executes any SQL statement.

The SQL statement will use the default database session, as will all methods of the DataSelection class.

```
begin method DoSelect  
empArray : Array of employee;  
sql select * into :empArray from emp; -- no on session clause  
end method;
```

► **To create a default database for use in a single method**

1. Assign the DBSession object to a local variable in your method:

```

etab : EmpObject = new;
DefaultDBSession : DBSession;
-- Assign to service object
DefaultDBSession = AnotherSession;
sql select * into :etab from tab; -- no on session clause
sql insert into tab values (:etab); -- no on session clause

```

Reconnecting to a Database Session

The Reconnect method of the DBSession class can be used to handle connections to a database that were initiated, but lost, by a DBSession service object. For example, you can use Reconnect to resume a connection dropped due to network failure or interruption.

The Reconnect method uses the same connection information that was used to make the previous connection, whether it was coded into the application, specified in the Service Object Properties dialog, or supplied through environment variables.

The value of the IsConnected attribute must be FALSE for the Reconnect method to execute. The occurrence of a DBRemoteAccessException sets the IsConnected attribute to FALSE.

A typical use of the Reconnect method is to invoke it upon the occurrence of a DBRemoteAccessException, as in the following example:

```

when e : DBRemoteAccessException do
    defaultDBSession.Reconnect();

```

It is most efficient to catch this exception and do the reconnect in the server partition (the same partition as the DBSession service object). You must do so if a router is in use, because otherwise the reconnect could be performed on a different replicate. If a router is not in use it is still preferred, since iPlanet UDS can catch the exception and perform the reconnect for one client, while other clients are unaffected by the interruption.

For more information on the Reconnect method, see the iPlanet UDS online Help.

NOTE Lost connections are treated differently for database sessions that are created dynamically by a DBResourceMgr service object. In this case, a TOOL service object can listen for the DBRemoteAccessException, and recreate a DBSession object in the exception handler.

Disconnecting a Database Session

Once the session is no longer needed, you end the session and release its resources by invoking the Disconnect method on the DBSession object, and setting the DBSession to NIL, as shown below:

```
MySession.Disconnect();  
MySession = NIL;
```

Setting the DBSession object to NIL is not required, but releases the object from memory immediately, which may be useful in more complex applications.

Vendor-Specific Notes

Informix

Informix has two mechanisms for connecting to Informix Dynamic Server databases: shared memory (local IPC) and network. A process, such as an iPlanet UDS server partition, is permitted only one shared memory connection. Multiple network connections are fully supported.

Oracle

iPlanet UDS uses OCI calls to communicate with Oracle. The connection is either a local Oracle connection or through SQL*Net V2.

SQL*Net V2 supports the following adapters: IPC and TCP/IP. IPC uses shared memory for local connections, and TCP uses the underlying transport protocol to connect to remote databases. Both IPC and TCP require the `tnsnames.ora` file.

► To make a connection using SQL*Net V2

- To make a local connection using the Bequeath adapter, use `@` as a database name.
- To make a local connection using the IPC adapter, use `@TNSResourceName` as a database name.

If you use the IPC adapter, the `sqlnet.ora` file must include the line:

```
automatic_ipc=on
```

- To make a network connection using the TCP adapter, use `@TNSResourceName` as a database name.

The `TNSResourceName` points to an entry in the `tnsnames.ora` file that SQL*Net V2 uses to identify the requested machine on behalf of the SQL*Net client; the entry specifies the adapter and Oracle SID, among other things.

If you have difficulty connecting to Oracle, try connecting using SQL*Plus to determine whether the problem is your Oracle configuration. A SQL*Net V2 connection might look like:

```
sqlplus joe/joepass@ORACLE_COACH
```

On Unix platforms you can see the shadow Oracle process once a connection is made and the command line will show what adapter is being used.

NOTE Oracle on OpenVMS has further architectural limitations. Oracle 7.3 on OpenVMS no longer supports shared memory database connections. In OpenVMS, regardless of the installed version of Oracle, an iPlanet UDS partition may have only one active “@” connection to Oracle. To create multiple connections, connect through the SQL*Net mailbox driver.

Rdb

You can connect to multiple Rdb databases in a single iPlanet UDS application by creating as many DBSession objects as you need. However, all concurrent DBSession objects created with a single DBResourceMgr service object must use the same user name. An application that needs multiple concurrent connections with different user names must use multiple DBSession or DBResourceMgr service objects.

DBSession or DBResourceMgr connections to Rdb use the “default” Rdb alias. If an application directly issues an Rdb ATTACH statement (to create, for example, multiple Rdb database attachments in a single Rdb transaction managed by DECdtm), it should specify an explicit, non-default, alias.

Working with Data Types

When you select data from a database table into an iPlanet UDS application, you must be aware of the type of data you are handling. iPlanet UDS converts database data to the iPlanet UDS data or class type that your application specifies. For many vendor data types, the appropriate iPlanet UDS types are fairly obvious, but some conversion issues are more subtle—for example, loss of precision or incompatible data types. You should choose iPlanet UDS data types carefully to make sure you represent the database data accurately and consistently.

This chapter discusses using simple (scalar) TOOL data types and more complex object data types with database data. It also contains several data type conversion tables, one for each database vendor, showing allowable mappings for when you insert iPlanet UDS data into a database, or select database data into iPlanet UDS objects or simple data types.

For reference information about simple TOOL data types, refer to the *TOOL Reference Guide*; for details about the object data types, refer to the online Help.

Using Database Data with iPlanet UDS

When you use database data in an iPlanet UDS application, one choice you have is whether to bring database data in as a simple scalar data type or as a class type. You also must indicate which scalar or class type is appropriate for each database value. Generally this means mapping database column data types to the appropriate iPlanet UDS data types. While some mappings are obvious and natural, there are more subtle considerations that you may very well encounter, including representing money, dates and times, internationalization issues, and numeric precision. This chapter addresses many of these issues; you may also wish to refer to the iPlanet UDS online Help or the *iPlanet UDS Programming Guide*.

In iPlanet UDS applications, you may store data in variables or in the attributes of objects, which are either scalar types or predefined subclasses of the iPlanet UDS `DataValue` class. In addition, some TOOL SQL extensions to some of the SQL statements allow for an automatic mapping of database result values to user-defined object attributes, and a mapping of user-defined object attributes to columns that will be inserted into a database.

Using Simple Data Types in TOOL

Simple TOOL data types are appropriate and easy to use when you transfer data between a database and an iPlanet UDS application and you do not encounter NULL values; for example, to retrieve integer data, you would simply declare a variable of type integer, and perform a select statement into that variable, as shown below:

```
year : integer;
sql select YearPainted into :year from Paintings
where Title = 'Mona Lisa'
on session MySession;
```

Using Nullable DataValue Subclasses

The Framework library contains several classes for data storage and manipulation, including subclasses to represent NULL values from database tables. The nullable classes are useful because database tables often contain NULL data. The following table lists the `DataValue` classes and corresponding nullable subclasses:

Class Data Type	Nullable Variant
BinaryData	BinaryNullable
BooleanData	BooleanNullable
DateTimeData	DateTimeNullable
DecimalData	DecimalNullable
DoubleData	DoubleNullable
ImageData	ImageNullable
IntegerData	IntegerNullable
IntervalData	IntervalNullable
TextData	TextNullable

You use these classes as you would any other iPlanet UDS class: define an attribute or variable of the class type you need, and then construct the object before you use it. For example:

```
Comments : TextData;
Comments = new;
```

Data Type Conversion

The next several sections contain tables that show the internal data types for DB2, Informix, ODBC, Oracle, Rdb, and Sybase databases, with corresponding TOOL data types. Note that the direction the data flows affects the appropriate data conversion. That is: certain TOOL data types can accept certain internal database types (for example, from a sql select statement), while the internal database types may accept other TOOL types (from a sql insert statement).

Handling NULL values Remember that if you select a NULL value into a scalar variable, scalar attribute, or non-nullable object, you will get an error of the `DatatypeException` class. To avoid this, use the nullable variant of the `DataValue` subclass (for example, `TextNullable`).

Overflow If you use a type that is too small, such as `i1` for a value bigger than 256, you may get conversion errors at runtime.

Reading the Data Type Conversion Tables

In the following data conversion tables, the first entry for each category represents the most “natural” mapping of the database type to an iPlanet UDS TOOL type. We have also grouped several related data types into one representative, generic TOOL data type. These groupings are shown in [Table 4-1](#).

Table 4-1 Legend for the Data Type Conversion Tables

Generic Data Type	Represents These Types
Boolean	boolean, BooleanData, BooleanNullable
Datetime	DateTimeData, DateTimeNullable
Integer	integer, long, short, i4, ui4, i2, ui2, i1, ui1, IntegerData, IntegerNullable
Interval	IntervalData, IntervalNullable
Text	String, TextData, TextNullable
Double	double, float, DoubleData, DoubleNullable, DecimalData, DecimalNullable
Binary	BinaryData, BinaryNullable, ImageData

For example, if an Oracle VARCHAR2 can be selected into a generic data type called “Text” it can also be selected into a String or objects of type TextData or TextNullable.

The following additional notes may apply to one or more of the data conversion tables.

- (a) Can be used as a host TOOL variable only if the selected column's data contains appropriate values for the TOOL data type (for example, a VARCHAR column can be selected into an IntegerData TOOL variable only if the column's value consists of numeric characters).
- (b) The database data type can contain values that are larger than can be represented in the TOOL variable data type. Selecting out-of-range values may result in undetected arithmetic overflow/underflow conditions (for example, a BIGINT can be selected into an IntegerData TOOL variable, but column values exceeding 2,147,483,647 in absolute value cannot be represented in an IntegerData object and will be incorrectly stored). Use these TOOL data types only if you are

certain that no out-of-range values will be stored in the database. Note that you may select these columns into a “large” TOOL data type (DoubleData, for example) and then assign it to another TOOL data type (for example, IntegerData) if its value is within IntegerData's range.

(c) When selecting a database column value that is “scaled” (for example, either a scaled integer or a floating point data type) into an exact numeric TOOL data type (IntegerData, i2, and so on), the fractional part of the value (the part to the right of the decimal point) will be lost.

(d) When inserting a TOOL data value into the database, the TOOL value must be appropriate for the database data type. For example, a BinaryData object inserted into a VARCHAR column cannot contain any embedded binary zeros, and a TextData object inserted into a TINYINT column must contain numeric characters and the total value must be within the -256 to +255 range for TINYINT.

(e) When inserting from a TOOL data type that has a larger range than the database data type (for example, from a TOOL DoubleData into a Rdb SMALLINT), overflow/underflow conditions may result. These may be undetected. Be careful when inserting from TOOL data types that are “larger” than their corresponding database data types. The iPlanet UDS application should ensure that out-of-range values are not inserted into the database.

(f) (Rdb and DB2 only) Rdb's DATE ANSI and DB2's DATE data types contain no hour/minute/second component. If the TOOL DateTimeData value contained such a value, that part will be lost. Use the Rdb or DB2 TIMESTAMP data type when values will include both year/month/day and hour/minute/second components.

(g) (Rdb and DB2 only) Rdb's and DB2's TIME data type contain no year/month/day component. If the TOOL IntervalData value contains a year/month/day component, it will be lost. Use the Rdb or DB2 TIMESTAMP data type when values will include both year/month/day and hour/minute/second components.

(h) (DB2 and ODBC only) Limitations in the DB2 CLI V2 interface prevent data type conversion for parameter markers. Parameter marker TOOL variables must therefore be of the data class that directly maps to the data type of the DB2 column or parameter marker. When using the DBSession classes directly (in lieu of TOOL SQL), it is possible to force data type conversion by explicitly setting the input marker's DBColumnDesc.DBDataType attribute to the value associated with the DB2 column's actual data type. For example, to insert an IntervalData object into a DB2 VARCHAR column, the input DataSet's associated DBColumnDesc.DBDataType attribute must be set to DB_DT_CHARACTERVARYING before the Execute method is invoked.

DB2 Data Conversion Table and Notes

In this release, TOOL parameter markers cannot be used for DB2 DATE data type columns. DateTimeData objects are assumed to be associated with DB2 TIMESTAMP columns. As a workaround, applications can affect DATE columns only by building queries using the DBSession classes directly. Specifically, when input values are set in the inputDataSet (returned by the Prepare method), the associated DBColumnDesc.DBDataType attribute must be changed to the value DB_DT_ANSIDATE. The Execute method will then be able to bind the parameter to the DB2 DATE column.

DB2's TIMESTAMP data type can contain fractional seconds to the resolution of a microsecond. iPlanet UDS DateTimeData data class, however, is limited to millisecond resolution. When DB2 TIMESTAMP columns are selected into iPlanet UDS DateTimeData objects, the fractional second component is truncated to millisecond resolution.

See notes (a) through (h) and [Table 4-1 on page 82](#) as you use this table.

Table 4-2 DB2 Data Conversion Table

DB2 Data Types	can be selected into these generic TOOL types	and inserted from these TOOL types
VARCHAR LONGVARCHAR CLOB	TEXT Integer (a) Boolean (a) DateTime (a) Double (a) Interval (a) Binary	TEXT (h)
CHAR	TEXT Integer (a) Boolean (a) DateTime (a) Double (a) Interval (a) Binary	TEXT (h)
GRAPHIC VARGRAPHIC LONGVARGRAPHIC	TEXT Binary	TEXT (h)
SMALLINT	INTEGER Double Text	INTEGER (e, h)

Table 4-2 DB2 Data Conversion Table (*Continued*)

DB2 Data Types	can be selected into these generic TOOL types	and inserted from these TOOL types
INTEGER	INTEGER	INTEGER (h)
DECIMAL	Double	
NUMERIC	Text	
DECIMAL	DOUBLE	DOUBLE (h)
FLOAT	Integer (b c) Text	
DATE	DATETIME Text	DATETIME (f) (see note above)
TIMESTAMP	DATETIME Text	DATETIME (h) (see note above)
TIME	INTERVAL Text	INTERVAL (g, h)
BLOB	BINARY Text	BINARY

Informix Data Conversion Table and Notes

The Informix DATETIME datatype is actually a family of 28 data types. Along with the DATE datatype, these are mapped to the iPlanet UDS DateTimeData class. When used in an SQL `where` clause, the `set` clause of an `update`, or as an input parameter to an Informix database procedure, only Informix DATETIME YEAR TO SECOND database columns and parameters can be referenced by a DateTimeData host variable. To reference columns of other DATETIME subtypes in these cases, use a TextData or string and specify the precise Informix string representation for the column or parameter.

The Informix INTERVAL data type is actually a family of 18 data types. These are mapped to the iPlanet UDS IntervalData class. When used in an SQL WHERE clause, the `set` clause of an `update`, or as an input parameter to an Informix database procedure, iPlanet UDS maps IntervalData values containing a year or month to an Informix INTERVAL YEAR TO MONTH value and other IntervalData values to Informix INTERVAL DAY TO FRACTION values. To reference other types of INTERVAL columns or parameters, use a TextData or string host variable and specify the precise Informix string representation for the column or parameter.

See notes (a) through (h) and [Table 4-1 on page 82](#) as you use this table.

Table 4-3 Informix Data Conversion Table

Informix Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
CHARACTER	TEXT	TEXT
CHAR	Integer (a)	Integer (d)
CHARACTER VARYING	Double (a)	Double (d)
NCHAR	DateTime (a)	Binary (d)
NVARCHAR	Interval (a)	DateTime (d)
TEXT		Boolean
VARCHAR		
INTEGER	INTEGER	INTEGER
INT	Double	Double (e)
SERIAL	Binary	Text (d)
SMALLINT	Boolean (a) Text	
DEC	DOUBLE	DOUBLE (e)
DECIMAL	Integer (c)	Integer (e)
DOUBLE PRECISION	Text	Text (d,e)
FLOAT	Boolean (a)	
MONEY		
NUMERIC		
REAL		
SMALLFLOAT		
BYTE	BINARY Text (a)	BINARY Text
DATE	DATETIME	DATETIME
DATETIME	Text	Text (d) (see note above)
INTERVAL	INTERVAL	INTERVAL (see note above)

ODBC Data Conversion Table and Notes

Only ODBC data types may be used with an iPlanet UDS ODBC resource manager. The mapping of database to ODBC data types is database-specific; consult your database driver documentation to obtain the mappings you require.

If you insert a data type of BIGINT with a fractional component into generic data type DOUBLE, it will be truncated (no rounding).

NOTE This manual does not include data type mappings for data sources such as SQL Server to ODBC, as those mappings are supported and updated by other vendors. For example, to see Microsoft SQL Server to ODBC datatype mappings you are referred to the documentation for Microsoft SQL Server.

See notes (a) through (h) and [Table 4-1 on page 82](#) as you use this table.

Table 4-4 ODBC Data Conversion Table

ODBC Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
VARCHAR LONGVARCHAR	TEXT Integer (a) Boolean (a) DateTime (a) Double (a) Interval (a) Binary	TEXT (h)
BINARY VARBINARY LONGVARBINARY	BINARY Text DateTime	BINARY
CHAR	TEXT Integer (a) Boolean (a) DateTime (a) Double (a) Interval (a) Binary	TEXT (h)
BIT	INTEGER Boolean Double	INTEGER Boolean Text (d)
SMALLINT	INTEGER Double Text	INTEGER (e) Double (e) Text (d)
TINYINT	INTEGER Double Text	INTEGER (e) Double (e) Text(d)

Table 4-4 ODBC Data Conversion Table (*Continued*)

ODBC Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
BIGINT	DOUBLE Integer (b) Text	DOUBLE (see note above) Integer Text (d)
INTEGER	INTEGER	INTEGER
DECIMAL	Double	Double
NUMERIC	Text	Text (d)
DECIMAL	DOUBLE	DOUBLE
FLOAT	Integer (b,c) Text	Integer Text (d)
DATE	DATETIME Text	DATETIME Text (d)
TIMESTAMP	BINARY	BINARY
TIME	DATETIME Text	DATETIME Text (d)

Oracle Data Conversion Table and Notes

An Oracle number type is integer if scale is 0, floating point otherwise.

See notes (a) through (h) and [Table 4-1 on page 82](#) as you use this table.

Table 4-5 Oracle Data Conversion Table

Oracle Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
VARCHAR2	TEXT	TEXT
CHAR	Integer (a)	Integer (d)
LONG	Double (a) DateTime (a) Interval (a)	Double (d) Binary (d) DateTime (d) Boolean
NUMBER	INTEGER (b) Double (see note above) Binary Boolean (a) Text	INTEGER Double (e) Text (d)
ROW ID	TEXT	Can't insert into ROWID

Table 4-5 Oracle Data Conversion Table (*Continued*)

Oracle Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
RAW MLS LABEL MLS LABEL	Not supported	Not supported
LONG RAW RAW	BINARY Text	BINARY Text
DATE	DATETIME Text	DATETIME Text (d)

Rdb Data Conversion Table and Notes

On the VAX only, Rdb's DOUBLE PRECISION columns are stored in "G_FLOATING" format, while TOOL's DoubleData data type is stored in "D_FLOATING" format. G_FLOATING's range is larger than D_FLOATING (see Digital's *VAX Architecture Handbook* for details), thus, very large column values (with exponents exceeding 38 in absolute value) cannot be stored in a TOOL data class. These cases are best handled by using CAST and/or arithmetic functions on the select statement to reduce the column value to a storable range. Note that on Alpha AXP OpenVMS, both Rdb DOUBLE PRECISION and TOOL DoubleData data types are stored in "G_FLOATING" format and are entirely compatible.

When inserting a TOOL data value that is "scaled" (either a DecimalData or DoubleData data type, for example) into an exact numeric Rdb column (TINYINT, SMALLINT, INTEGER, or BIGINT), the fractional part of the value (the part to the right of the decimal point) will be lost.

When inserting into a LIST OF BYTE VARYING column, the byte string is written in segments of the maximum segment size defined when the column was originally created. When these columns are fetched, all the columns' segments are concatenated to form the BinaryData object. Note that Rdb SQL does not permit LIST OF BYTE VARYING columns to be updated; they may only be inserted and selected.

DBKEYS cannot be inserted into a table, but can be used as host variables in a WHERE clause to identify a row to be updated, deleted, or fetched.

See notes (a) through (h) and [Table 4-1 on page 82](#) as you use this table.

Table 4-6 Rdb Data Conversion Table

Rdb Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
VARCHAR	TEXT Integer (a) Boolean (a) DateTime (a) Double (a) Interval (a) Binary	TEXT Integer Boolean DateTime Double Interval Binary (d)
CHAR	TEXT Integer (a) Boolean (a) DateTime (a) Double (a) Interval (a) Binary	TEXT Integer (e) Boolean (e) DateTime (e) Double (e) Interval (e) Binary (d,e)
TINYINT (unscaled)	INTEGER Double Text	INTEGER (d,e) Double (d,e) Text (d,e)
TINYINT (scaled)	DOUBLE Integer (c) Text	DOUBLE (d,e) Integer (d,e) Text (d,e)
SMALLINT (unscaled)	INTEGER Double Text	INTEGER (d,e) Double Text (de) (see note above)
SMALLINT (scaled)	DOUBLE Integer (c) Text	DOUBLE (e) Integer (e) Text (d,e)
INTEGER (unscaled)	INTEGER Double Text	INTEGER Double (e) Text (d,e) (see note above)
INTEGER (scaled)	DOUBLE Integer (c) Text	DOUBLE (e) Integer (e) Text (d,e)

Table 4-6 Rdb Data Conversion Table (*Continued*)

Rdb Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
BIGINT (unscaled)	DOUBLE Integer (b) Text	DOUBLE Integer Text (d,e)
BIGINT (scaled)	DOUBLE Integer (b,c) Text	DOUBLE (e) Integer Text (d,e)
REAL	DOUBLE Integer (b c) Text	DOUBLE Integer Text (d)
DOUBLE PRECISION	DOUBLE (see note above) Integer (b,c) Text	DOUBLE Integer Text (d)
DATE VMS	DATETIME Text	DATETIME Text (d)
DATE ANSI	DATETIME Text	DATETIME (f) Text (d,f)
TIMESTAMP	DATETIME Text	DATETIME Text (d)
TIME	INTERVAL Text	INTERVAL (g) Text (d,g)
INTERVAL	INTERVAL Text	INTERVAL Text (d)
LIST OF BYTE VARYING (see note above)	BINARY Text (a)	BINARY Text (see note above)
DBKEY	BINARY	BINARY

Sybase Data Conversion Table and Notes

See notes (a) through (h) and [Table 4-1 on page 82](#) as you use this table.

Table 4-7 Sybase Data Conversion Table

Sybase Data Type	can be selected into these generic TOOL types	and inserted from these TOOL types
VARCHAR	TEXT	TEXT
CHAR	Integer (a) Boolean (a) Double (a) DateTime (a) Binary Interval (a)	Binary (d)
TINYINT	INTEGER	INTEGER (d)
SMALLINT	Double	Double (e)
INT	Text Boolean	Boolean
REAL	DOUBLE	DOUBLE
FLOAT	Integer (b)	Integer
MONEY	Text	Boolean
SMALLMONEY	Boolean	
DATETIME	DATETIME	DATETIME
SMALLDATETIME	Text	Text
VARBINARY	BINARY	BINARY
BINARY	Text DateTime	Integer Boolean
BIT	INTEGER Text Boolean	INTEGER Boolean
TEXT	TEXT Binary	TEXT Binary DateTime
IMAGE	BINARY Text	BINARY Text DateTime

Manipulating Data

The iPlanet UDS database interface allows application programmers to write applications that use the full potential of relational databases. This chapter describes how to add SQL statements to an iPlanet UDS application so that it can access and manipulate database data.

You can use TOOL SQL statements to query or update data. You can also use the iPlanet UDS database management classes found in the GenericDBMS library to access and manipulate database data.

This chapter describes the various ways you can query, insert, update and delete database data. It includes several examples, some drawn from the iPlanet UDS examples. It also describes creating and naming the iPlanet UDS objects into which you will select data. At the end of the chapter are database vendor-specific notes.

For reference information about TOOL SQL statements, refer to [Appendix B, "TOOL SQL Statement Reference."](#)

Accessing Database Data from iPlanet UDS

While iPlanet UDS provides full support for building SQL statements into an iPlanet UDS application, programmers should keep in mind some general questions when working with database data. These questions include:

- If I use non-standard SQL (that is, vendor-proprietary extensions to SQL) in my application, will that prevent my application from working with any database that it must work with?
- When should I use TOOL SQL statements or the classes in the GenericDBMS library?
- When moving data between a database and iPlanet UDS application, how can I assure that I use compatible data types?

iPlanet UDS passes all SQL to the database including vendor-specific extensions. Using only standard ANSI SQL increases the likelihood that your applications will run on any of the database management systems that iPlanet UDS supports. Note, however, that iPlanet UDS does not do any translation of TOOL SQL before passing it to the DBMS; that is, compiling a TOOL method containing SQL does not guarantee that it will execute on any database. If you need to write an application that is portable across multiple databases, you must ensure that you use a subset of SQL that works for all of those particular databases.

A small number of vendor-specific SQL extensions are not supported; to see a list, refer to [“Unsupported Database Features” on page 35](#).

To manipulate database data, you can use either TOOL SQL statements or iPlanet UDS GenericDBMS classes.

TOOL SQL statements are more convenient to use, particularly when you can embed specific SQL statements in your application. An example is a query for which you know the selection criteria and the names of the columns to be returned, even though you might not know what the selection criteria *values* are. You can embed the known column and table names, and use variables to get a selection value from the user (such as, generate the report for Department “10”).

However, TOOL SQL cannot be used to write SQL statements that are generated at runtime. An application that allows SQL statements to be generated and executed “on-the-fly” has the potential to be a more flexible, multi-purpose application. To write such an application, you use the classes in the iPlanet UDS GenericDBMS library, particularly the DBSession class. The section [“Using iPlanet UDS Classes to Execute SQL” on page 118](#) begins the discussion of using the DBSession class interface to write these applications.

Application performance is the same, whether you use TOOL SQL or methods on the DBSession class, because TOOL statements are translated into invocations of methods on the DBSession class at compile time.

In iPlanet UDS applications, you may store data in variables or in object attributes, which are either scalar types or predefined subclasses of the iPlanet UDS DataValue class. In addition, TOOL SQL extensions to some of the SQL statements allow for an automatic mapping of database result values to user-defined object attributes, and a mapping of user-defined object attributes to columns that will be inserted into a database.

You must also consider issues of data type conversion when integrating database data into your iPlanet UDS application. While examples in this chapter illustrate some data type conversion, you should also refer to [Chapter 4, “Working with Data Types”](#) to see how iPlanet UDS converts and uses the data types for your particular database vendor.

Using iPlanet UDS Names in SQL Statements

When you use iPlanet UDS names (variables, attributes, and so on) in TOOL SQL statements, you must preface the iPlanet UDS names with colons to distinguish them from database names (columns and tables).

Code Example 5-1 Distinguishing iPlanet UDS names from database names in TOOL SQL statements

```
name : TextData = new;
year_born : integer = 1900;
sql select ptr_name into :name from painter_table
  where birth < :year_born;
```

NOTE You do not precede iPlanet UDS names with colons in two TOOL SQL statements (sql open cursor and sql execute procedure). These statements both take a parameter list and a syntax error occurs if a colon precedes a parameter value.

TOOL SQL Statements

The following table lists the TOOL statements that correspond to SQL DML statements used to query and update database data. You can find reference information for these statements in [Appendix B, “TOOL SQL Statement Reference.”](#)

TOOL SQL Statement	Description
sql select	Retrieves rows from a database table.
sql delete	Removes rows from a database table.
sql insert	Adds a new row to a database table.
sql update	Replaces values in a database table.
sql open cursor	Opens a cursor.
sql close cursor	Closes a cursor.
sql fetch cursor	Retrieves rows from a cursor.

TOOL SQL Statement	Description
sql execute immediate	Executes a single SQL statement specified as a literal string, a string variable, or a TextData variable.
sql execute procedure	Executes a database procedure.

For a TOOL SQL statement to execute in a database, it must be associated with a connected database session. You can initiate a database connection in three ways, described in [“Connecting to a Database” on page 55](#). To associate a TOOL SQL statement with a database session, you use the `on session` clause with the name of the DBSession object that you wish to use for the connection. Note that not all statements require or allow the `on session` clause.

For example, the following query retrieves a single row from the database, using the session specified by the MySession DBSession object, into an object referenced by a variable called emp:

```
sql select * into :emp from emptable
where name = 'Smith'
on session MySession;
```

Using Conditional TOOL for Vendor-Specific Code

In your SQL statements, you can use any vendor-specific extensions to the ANSI SQL syntax that are allowed by the particular database management system you are using. Examples of extensions to ANSI standard SQL include the Rdb cast statement, Sybase compute statement, or Oracle outer join.

However, if you include a vendor-specific clause in a SQL statement, your code is no longer generic, and will work only for database products that support that particular syntax. Also note that because iPlanet UDS passes vendor-specific syntax directly to the database, it does not detect syntax errors within vendor-specific clauses.

If your application must execute against multiple databases and you want to use vendor-specific SQL extensions, then you can use conditional TOOL code. Conditional code allows you to issue different SQL statements depending on which type of database an application is connected to.

To write conditional code that is based on the database vendor, you use the `DBVendorType` attribute of the `DBSession` class. This attribute contains the database vendor for the current database session. By using the `DBVendorType` attribute in a case or if statement, you can include vendor-specific SQL extensions that will be invoked only for the databases that support them.

The following example shows the use of `DBVendorType` to issue vendor-specific SQL:

```
case DefaultDBSession.DBVendorType
when DB_VT_ORACLE do
  sql execute immediate
    'create table emp (id int, hired date, salary float)';

when DB_VT_INFORMIX do
  sql execute immediate
    'create table emp (id integer, hired datetime year to fraction
(3),
  salary money)';

when DB_VT_SYBASE do
  sql execute immediate
    'create table emp(id int, hired dateTime, salary float)';

when DB_VT_DB2 do
  sql execute immediate
    'create table emp (id integer, hired timestamp, salary
float)';

when DB_VT_Rdb do
  sql execute immediate
    'create table emp (id integer, hired timestamp, salary
double precision)';
end case;
```

Using TOOL Statements to Query Data

The following sections describe various ways of retrieving data from a database, with several examples. Also note the first few sections that describe how to create and name objects and variables into which you will select database data.

Selecting Data and Object Creation

When you query data (using either the `sql select` or `cursor` statements) you may assign the retrieved data to an iPlanet UDS object. In some cases, iPlanet UDS creates the object for you; in other cases you must explicitly create the objects yourself.

In a simple select into a set of `DataValue` objects, you must create the `DataValue` objects before the select, as in the following:

```
i : IntegerNullable = new;
f : DoubleNullable = new;
sql select intval, floatval into :i, :f
  from myTab on session MySession;
```

You must also create objects that are used as attributes to the object. For example:

```
class myClass
  has public i : IntegerNullable;
  has public f : DoubleNullable;
end class

method init
  begin
    super.Init;
    self.i = new;
    self.f = new;
  end method;

method doSelect
  begin
    myObj : myClass = new;
    sql select * from myTab into :myObj on session MySession;
```

You do not need to create an instance of the array when you select into an array of objects. iPlanet UDS also creates the rows of the array during the select. For example:

```
myArray : Array of myClass;
sql select * from myTab into :myArray on session MySession;
```

iPlanet UDS automatically creates the objects on each iteration in a for loop, as shown below:

```
for (myObj : myClass) in sql select * from myTab
  on session mySession do
  ... myObj is created for you...
end for;
```

You must explicitly create the objects before a `sql fetch cursor` statement, as shown below:

```
-- Assume cursor myCursor with the following definition
-- select * from myTab on session mySession

myCursor_ref : myCursor;
sql open cursor myCursor_ref on session mySession;
while (TRUE) do
  i : integer;
  i : IntegerNullable = new;
  f : DoubleNullable = new;

  i = sql fetch cursor myCursor_ref into :i, :f;
end while;
```

Selecting a Single Row

You can select a single row from a database into an iPlanet UDS variable or object attribute using the TOOL sql select statement. If more than one row would be returned, iPlanet UDS raises an exception of the DBOperationException class, and does not return any data. To select more than one row, you must select into an array or a cursor, or use a for loop—these techniques are explained below.

Selecting into a Variable

You can select data from a database into one or more iPlanet UDS variables. In the following example, note that the two iPlanet UDS variable names are preceded with colons when used in the TOOL statement:

```
name : string = 'Picasso';
t_comments : TextData = new;
t_country : string;
sql select comments, country into :t_comments, :t_country
from ArtistTab
where Name = :name
on session MySession;
```

Selecting into an Object

You can use the sql select statement to retrieve one row from a database table and store the row values in an iPlanet UDS object. To do so, the names of the object attributes must match the names of the retrieved columns. (Attribute names are case-insensitive, while column names may be case-sensitive depending upon the database vendor.) You specify the conditions that the row must meet, as in the following example:

```
painter : Artist = new;
name : TextData = new(value='De%');
sql select * into :painter
from ArtistTab
where Name like :name
on session MySession;
```

When Attribute and Column Names Match

When a `sql select` statement selects all columns from a table into an object (for example, using `select *`), iPlanet UDS retrieves all table columns. Every column with a matching attribute is selected into that attribute. That is, if the column name and attribute name match (case-insensitive), the attribute receives the column's value.

If the table has columns with no corresponding attribute, those columns are ignored. Similarly, if the object has attributes with no corresponding columns, those attributes are untouched. Note that if the number of attributes is significantly higher than the number of columns, you will get better performance if you explicitly specify the column names in the select statement.

If a `sql select` statement selects specific columns from a table, those columns must have matching attributes in the receiving object. For example, if you select the name and birthdate from `ArtistTab`, then the painter object must also have attributes named name and birthdate.

When Attribute and Column Names Do Not Match

When you need to select into an object whose attribute names do not match the column names of the source database table, you can do so in several ways:

You can “rename” the columns in the select list of the statement, as in:

```
sql select col1 "attr1", col2 "attr2", into :obj from tablea...;
```

You can select directly into the attributes, as in:

```
sql select col1, col2 into :obj.attr1, :obj.attr2 from tableA...;
```

You can also use the `TOOL for` statement to achieve column renaming in order to match column and attribute names, as described in the next section.

When Attributes are of Other Class Types

In most cases, if there are any attributes in the object that have user-defined class types (other than `DataValue` types), iPlanet UDS ignores these attributes even when they have matching names. To select into such an object, you must rename the column in the select list to reference the class to which the attribute belongs. For example, assume you have the two classes `AddrClass` and `EmpClass`:

```
class AddrClass inherits from Object
has public
  Street : TextData;
  State : TextData;
  Zip : Integer;
  ...
```

```
Class EmpClass inherits from Object
has public
  Name : TextData;
  Address : AddrClass;
```

Also assume the `Emp` table has the following definition:

```
table Emp
(ID int, Name char(20), Street char(30), City char(20), Zip int)
```

To select into an `EmpClass` object, which contains an attribute `Address` of type `AddrClass`, you must rename the columns that reference the `Street`, `City`, and `Zip` attributes, as shown below:

Code Example 5-2 Renaming Columns

```
empObj: EmpClass;
sql select
  ID,
  Name,
  Street      "Address.Street",
  City        "Address.City",
  Zip         "Address.Zip"
into :empObj
from Emp on session MySession;
```

NOTE Column renaming may not be portable across all databases and is subject to restrictions of each database vendor. Some of the known vendor-specific restrictions follow:

- For `Rdb`, columns are renamed using the ANSI SQL “as” clause.
- For `Oracle`, the length of a column alias is limited to 32 characters.
- `DB2` does not allow you to rename a result column arbitrarily (eliminating the above approach).
- `Informix` does not allow an embedded period when renaming a column (also eliminating the above approach).

If the syntax in the example above does not work for your database, you can use a `FOR` loop to achieve the same effect; see the *TOOL Reference Guide* for more information on the `FOR` statement.

When Attributes are Inherited

As is true with all TOOL objects, you can select data into an object that is a subclass of another object, using the attributes of both classes. For example, say you have a superclass `Person` with attributes `Name`, `Address`, and `Age`. `Person` has a subclass called `Employee`, with the attribute `Salary`. Your database table has columns `Name`, `Address`, `Age`, and `Salary`. You can successfully perform the following select statement:

```
info : Employee = new;
sql select Name, Address, Age, Salary into :info
    from empTable on session MySession;
```

Selecting Multiple Rows into Arrays

To select data into an array, you declare an array object and select into it (you do not need to instantiate it). Also, you do not need to add rows to the array—when you perform the select, the array will accommodate as many rows as are retrieved. Any rows that existed in the array before the select will be gone.

You can use the `into` clause to select multiple rows into an array, as shown in the following example:

```
painter.ListOfPaintings : Array of Painting;
sql select * into :painter.ListOfPaintings from ArtistTab
    where Painter = :painter.name
    on session MySession;
```

Selecting Multiple Rows Using the TOOL for Statement

You can use the `for` statement with a `sql select` statement to retrieve multiple rows from a database.

```

painters : Array of Artist;
for (a : Artist) in sql select * from artistTab
on session MySession do
-- A new Artist object is allocated each time through
  painters.AppendRow(a);
end for;

```

Selecting Multiple Rows using Cursors

A cursor is a row marker that you use to work with a set of rows from a database. Cursors have the following benefits:

- Cursors are useful if you do not know how many rows a query will retrieve—you can simply fetch one or more rows at a time from the cursor's result set.
- A cursor is defined in one place only—the Cursor Workshop. You can use a cursor in multiple methods, but if you need to modify the cursor's select statement, you need only make the change once in the Cursor Workshop.

You can use a cursor in several different ways:

- You can use the `sql open cursor`, `sql fetch cursor`, and `sql close cursor` statements to step through the cursor's result set one row at a time.
- You can use the same statements to fetch the entire result set into an array.
- You can use the same statements, but with `sql fetch next n`, to step through the result set n rows at a time.
- You can use the `for` statement to repeat a statement block for each row in the result set of the cursor.

For Oracle, Rdb, and Informix, you can also use a cursor with the `sql update ...` where `current of or sql delete ...` where `current of` statements to update or delete the row to which the cursor is pointing (a “positioned update”). When you define the cursor, you need to include the `for update` clause in the cursor’s SQL statement. See [“SQL Update” on page 190](#) and [“SQL Delete” on page 169](#) for more information about using positioned updates.

Defining a Cursor

You define a cursor in the Cursor Workshop (described in *A Guide to the iPlanet UDS Workshops*). A cursor definition includes a cursor name and an underlying query. Because a cursor definition is associated with an iPlanet UDS project, you can refer to the cursor in any method in the project.

An example of a cursor definition follows:

```
cursor BlobCursor(name: Framework.string)
begin
  select BlobValue from ArtistBlob
  where Name like :name;
end;
```

As described in the *TOOL Reference Guide*, most TOOL statements use a cursor reference rather than the cursor’s name. (The TOOL `for` statement uses the cursor name rather than a cursor reference.) In the following code fragment `blobCurs` is the cursor reference for the cursor named `BlobCursor`:

```
blobCurs : BlobCursor;
begin transaction
  sql open cursor blobCurs(name) on session self.MGRSession;
  rowcount = (sql fetch cursor blobCurs into :binData);
  sql close cursor blobCurs;
. . .
```

Retrieving Rows

To retrieve rows from a cursor, you must first open the cursor with the `sql open cursor` statement. You can then use the `fetch` statement to retrieve one or more rows at a time. Finish by using the `sql close cursor` to close the cursor. For best performance, you should enclose all statements relating to one use of a cursor in an explicit iPlanet UDS transaction.

The `sql open cursor` statement executes the `select` statement associated with the cursor and positions the cursor before the first row in the result set. At this point you specify the values for any placeholders used in the original cursor declaration. (Placeholders are mechanisms through which you can pass data into your methods at runtime.):

```
dbcursor : empcursor;
empid : integer;
sql open cursor dbcursor (empid) on session MySession;
```

In this example, the cursor name is `empcursor` and the cursor reference is `dbcursor`.

Use the `sql fetch cursor` statement to move one or more rows at a time through the result set. This statement need not be in the same method as the `sql open cursor` statement but it should be in the same explicit transaction. The first time you use the `sql fetch cursor` statement, iPlanet UDS moves the cursor to the first row (or set of rows) in the result set and retrieves the data into the specified TOOL variables. With each successive `sql fetch` statement, iPlanet UDS moves the cursor forward one or more rows, retrieving the data into the specified TOOL variables. You continue using `sql fetch cursor` to move through the result set until you reach the last row.

```
emp : employee = new;
sql fetch cursor dbcursor into :emp;
```

Use the `sql close cursor` statement to close the cursor. After the cursor is closed, it cannot be used again until you give another `sql open cursor` statement to open it.

```
sql close cursor dbcursor;
```

Fetching into an Array

You can also fetch the entire result set into an array. For example:

```
empArray : array of employee;
sql fetch cursor dbcursor into :empArray on session MySession;
```

Fetching an Arbitrary Number of Rows

Rather than having to fetch only one row or all rows, you can use the `sql fetch cursor` statement with the `next` key word to specify the maximum number of rows to fetch into an array. For example:

```
dbsess : DBSession;
rows   : integer = 1;
dbcursor : artist_cursor;
painters : array of artist = new;

sql open cursor dbcursor ('%') ON SESSION dbsess;
while (rows > 0) do
    sql fetch next 10 from cursor dbcursor into :painters;
    rows = painters.Items;
    -- do something with painters ...
end while;
```

Repeating a Statement Block

You can use a `for` statement to create a loop that will fetch and process one record at a time. The `for` statement automatically opens the cursor, fetches the rows one at a time as it goes through the loop, and then closes the cursor. To use the `for` statement with a cursor, you use the cursor name rather than a cursor reference, as follows:

```
MySession : DBSession = ...get a session...
for (emp : employee)
    in cursor empcursor(empid) on session MySession do
        -- A new emp object is allocated each time through
        -- user code
end for;
```

Note that this operation is identical to using the `for` statement with a `sql select` statement (described above). Using a cursor in the `for` statement is especially useful when the cursor is also used in other parts of the application.

Using TOOL to Update Data

The next several sections discuss using TOOL SQL statements to insert, update, and delete database data.

Inserting a Single Row

You use a `sql insert` statement to insert data directly into a table.

Inserting Variables

You can insert data from variables using the `sql insert` statement. For example:

```
t_comments : TextData = new(value = 'Comments');
t_name : string = 'Picasso';
t_country : string = 'Spain';
sql insert into ArtistTab
    (name, comments, country)
    values (:t_name, :t_comments, :t_country)
    on session MySession;
```

Inserting from an Object

Use the `sql insert` statement to add a new row to a table from an object. If you do not specify a column list, iPlanet UDS uses the object's attributes as the column list. If you do not specify a column list and the database uses case-sensitive column names (for example, Sybase or Microsoft SQL Server), the attribute names must exactly match the database column names (that is, both spelling and case must match).

```
NewPainting : Artist = new;
sql insert into ArtistTab
    values (:Newpainting)
    on session MySession;
```

You can also insert into a subset of the columns by specifying a column list:

```
NewPainting : Artist = new;
sql insert into ArtistTab(name, comments)
  values (:newPainting)
  on session MySession;
```

NOTE If you do not instantiate an object (using “new”), it is treated as a NULL by the database.

Inserting Multiple Rows

To insert a set of rows into a table, you can use a `sql insert` statement with an array in the `values` clause. This is the most efficient way to insert a set of rows. For example:

```
artists : Array of Artist = new;
-- Fill in array...
sql insert into ArtistTab values (:artists) on session MySession;
```

Updating a Row

To replace the current column values in selected rows, use the `sql update` statement. For example:

```
sql update ArtistTab set born = :birthyear
  where name = :vname on session MySession;
```

Unless you specify a `where` clause, values in all rows are updated.

Deleting a Row

To remove one or more rows from a table, use the `sql delete` statement. For example:

```
sql delete from ArtistTab where born < 1500;
```

If you do not include a `where` clause, all rows in the table are deleted.

Executing a Single SQL Statement

Use the `sql execute immediate` statement to execute a SQL statement that you enter directly as a literal string, or a statement that is stored in an iPlanet UDS variable or attribute. You use the `sql execute immediate` statement to execute SQL statements that are not explicitly part of TOOL, or to execute statements constructed at runtime. For example, you can use it to issue DDL (data definition language) statements, such as the following:

```
sql execute immediate 'create table ArtistTab(
  Name char(30) not null, Country char(30) not null,
  Comments varchar(200) not null)'
  on session MySession;
sql execute immediate
  'grant all on ArtistTab to public' on session MySession
```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** MakeTables

Executing a Database Procedure

To execute a database procedure, use the `sql execute procedure` statement. You can pass parameters either by name or by position. For example:

```
empid : integer = 12345;
salaryIncrement : integer = 15000;

-- Passing parameters by position.
sql execute procedure updateSalary(empid, salaryIncrement);

-- Passing parameters by name.
sql execute procedure updateSalary(AddToSalary = salaryIncrement,
  Id = empid);
```

Vendor-Specific Notes on Database Procedures

The following table describes iPlanet UDS support for database procedures:

DBMS	Procedures	Named Parameters	Positional Parameters	Comments
DB2 v2	Yes	No	Yes	v2 CLI ODBC restriction
Informix	Yes	Yes	Yes	no input/output parameters
ODBC	Yes	No	Yes	drivers may provide less, eg Access
Oracle	Yes	Yes	Yes	
Rdb	Yes	No	Yes	
SQLserver	Yes	No	Yes	ODBC restriction
Sybase	Yes	Yes	Yes	

- Informix dbprocs do not support input-output parameters. If input-output parameters are passed from a TOOL execute procedure, an exception is generated.
- If the ODBC database driver you are using does not support procedures, you will receive a runtime error.

Oracle procedures can return functional results. The following example is valid for Oracle:

```
intval : integer;
intval = IntegerData((sql execute procedure p(param1 =
5))).Value;
```

Because `sql execute procedure` returns a `DataValue` object, you must cast the return value to the desired class. In this example, an extra set of parameters is required in order to cast the result to `IntegerData`. Thus, you must know what datatype the procedure returns so as to cast it appropriately. If the cast is to the wrong class, the user receives a runtime error.

Working with ImageData Objects

When selecting ImageData from a database, the data is read using one of several iPlanet UDS formats, just as is described for the ReadFromFile method for the ImageData class. iPlanet UDS determines the type of image just as it does when the format SP_IF_DEFAULT is used with ReadFromFile; that is, iPlanet UDS determines the type of image data by reading the start of the actual data.

When inserting ImageData into a database, the data is always written as serialized ImageData.

Using Binary Large Objects (BLOBs)

You can generally manipulate BinaryData objects (also called Binary Large Objects, or BLOBs) and large TextData objects just like other datatypes. The iPlanet UDS database interface allows you to fetch, insert, and update BLOB columns (both TEXT and BYTE) using the BinaryData and TextData classes.

You need not know in advance the maximum BLOB size to select BLOB data. iPlanet UDS automatically allocates enough space to retrieve or manipulate any BLOB data.

When working with binary objects, you must be sure to satisfy vendor-specific requirements. To ensure application portability, follow the recommendations below when working with binary data.

Selecting Binary Data

The following example demonstrates using a cursor to fetch BLOB from a database into a BinaryData object. It then deserializes the BinaryData object into an object of type Artist. The cursor reference BlobCurs refers to the cursor that is defined in [“Defining a Cursor” on page 106](#).

```
GetArtist(name:string) : Artist

-- This method returns an Artist Object, given an artist
-- name. It does this by fetching a serialized version
-- of the Artist from the ArtistBlob table in the database.

-- When it gets the BinaryData, it is assumed to contain
-- a MemoryStream object which is deserialized.
```

```

-- The method returns an Artist object, or NIL value, if no
-- artist with matching name is found.

-- This will pass any exceptions back to the invoking method
-- for display, including an exception if more than one
-- matching name was found.

-- First, get the BinaryData if it exists.
rowcount : integer = 1;
binData : BinaryData = new;
blobCurs : BlobCursor;

begin transaction

    sql open cursor blobCurs(name) on session self.MGRSession;
    rowcount = (sql fetch cursor blobCurs into :binData);
    sql close cursor blobCurs;
    if rowcount = 0 then
        return NIL;
    end if;

end transaction;

-- Now that the data is read from the DB, deserialize
-- the MemoryStream that is in the BinaryData.
strm : MemoryStream = new;
strm.Open(accessMode = SP_AM_READ, isBinary = TRUE);
-- The UseData method will simply make the stream point
-- to the memory that has already been set up in the
-- BinaryData object.
strm.UseData(data = (pointer to char)(binData.Value),
    length = binData.ActualSize);
-- Deserialize the Artist. Remember to cast.
painter : Artist =
    (Artist) (strm.ReadSerialized());
strm.Close();

return painter;

Project: WinDB • Class: ArtistMgrTable • Method: GetArtist

```

Inserting Binary Data

When inserting BLOB data into a database, iPlanet UDS generates multiple SQL statements from the original `sql insert` statement, and, in turn, requires information about column names that correspond to the data inserted. Thus, to insert a BLOB, you must use an object of a class whose attribute names and data types correspond to the columns you wish to insert into a table.

The following code example shows serializing an iPlanet UDS object into a BinaryData object, and then inserting the BinaryData object into a database column declared as a BLOB type.

```

MakeDatabase()

-- This method creates the blob table needed in the example.
-- A single table is created, with an appropriate type for
-- the target database (since each supports a different
-- type for blobs). It is created with two columns,
-- one for the key (artist name) and one with the actual blob
-- data. The sample data is loaded first into memory,
-- and then transferred to the blob.

-- Any exceptions that occur will be passed back to the
-- code that invokes this method. This allows the user
-- interface to display any of the error code.

-- First, get some sample data, from a generic method.
artists : Array of Artist;
artists = self.GetSampleData();

-- The tables are assumed not to exist. Any errors if
-- they do are simply sent back to the client. Note that
-- the names of the columns use the exact same case as the
-- attributes in the tables. This is because Sybase
-- is case sensitive in names. Also the table name
-- is case sensitive.

case self.MGRSession.DBVendorType is

when DB_VT_SYBASE do
    sql execute immediate
        'create table ArtistBlob(Name varchar(30) not null,
BlobValue image not null)'
    on session self.MGRSession;
when DB_VT_ORACLE do
    sql execute immediate
        'create table ArtistBlob(Name varchar(30) not null,
BlobValue long raw not null)'
    on session self.MGRSession;
else do
    e : GenericException = new;
    e.SetWithParams(severity = SP_ER_ERROR,
        message = 'This example does not demonstrate using blobs
for this database.');
```

```

        task.ErrorMgr.AddError(e);
        raise e;
end case;

-- Grant privileges
sql execute immediate
    'grant all on ArtistBlob to public' on session self.MGRSession;
```

```

-- Create memory streams for the artist data, transfer them into
-- BinaryData objects, and then insert into the database.

-- Note also that the data is transferred into the
artistBlobObjects
-- array. The row class of this array has attributes that
-- correspond to the columns in the ArtistBlob table. This is
-- a more portable way to handle insertion of blob data.

artistBlobObjects : array of ArtistBlobObject = new;
i : integer = 1;
while i <= artists.Items do
    strm : MemoryStream = new;
    strm.Open(accessMode = SP_AM_READ_WRITE, isBinary = TRUE);
    strm.WriteSerialized(object = artists[i]);
    -- Now transfer the stream to a BinaryData object.
    strm.Seek(0);
    artistBlobObjects[i] = new;
    artistBlobObjects[i].Name = artists[i].Name;
    strm.ReadBinary(target = artistBlobObjects[i].BlobValue);

    -- Insert into the database
    sql insert into ArtistBlob(Name, BlobValue)
        values (:artistBlobObjects[i])
        on session MGRSession;
    strm.Close();
    i = i + 1;
end while;

return;

```

Project: WinDB • **Class:** ArtistMgrBlob • **Method:** MakeDatabase

Vendor-Specific Notes on BLOB Handling

iPlanet UDS does not support references to Informix BLOBs by file name or file handle.

Because Rdb does not allow updates of BLOB data, you cannot update Rdb BLOB data using either TOOL or the iPlanet UDS classes. However, you can insert BLOB data.

There are some additional restrictions when inserting or updating BLOB data with Sybase. These restrictions are necessary since the Sybase interface does not allow direct modification of Sybase text or image columns through SQL. Text or image columns must be modified by selecting back the affected row and then using a separate interface.

The restrictions on using BLOB data with Sybase follow:

1. If you insert a row which contains a text or image column, then all of the following must be true:
 - a. The non-text, and non-image columns must uniquely identify the row to be inserted. In the following example the variables `intval` and `charval` must contain values that together will uniquely identify the row inserted:

```
SQL INSERT INTO T (intcol, charcol, textcol) values
(:intval, :charval, :textval);
```

- b. You must explicitly name the columns to be inserted.

Unlike the previous example, the following example will fail since it does not have a column list:

```
SQL INSERT INTO T values (:intval, :charval, :textval);
```

This restriction does not apply if you do not insert any text or image columns.

- c. The `values` clause cannot contain literals. For example, the following example will fail because the literal `5` is invalid in this context:

```
SQL INSERT INTO T (intcol, charcol, textcol) values
(5, :charval, :textval);
```

2. If you insert or update a row with more than one text or image column, then *all* the text and image values must be either less than 256 bytes, or 256 bytes or more. If some values are shorter and some longer then you must break the statement into two, one which inserts or updates the text and image columns shorter than 256 bytes, and another which updates the values 256 bytes or longer.
3. When performing an update which contains a text or image column, you must include a `where` clause which limits the update to a single row.
4. When performing an update which contains a text or image column you cannot update any column which appears in the `where` clause.

Using iPlanet UDS Classes to Execute SQL

Whenever possible, it is most convenient to use TOOL SQL statements directly in your methods. However, for many applications, you will need to use the GenericDBMS database class interface. This interface lets you execute a SQL command whose exact text is unknown at compile time; it provides more flexibility, but may require more programming to handle various conditions.

You can generate SQL statements at runtime in your application when any of the following is unknown:

- the text of the SQL statement
- the number and/or data types of host variables (either for input or what will be returned by the SQL statement)
- the references to columns, tables, and views

iPlanet UDS sample applications Two iPlanet UDS example applications, DynamicSQL and DynamicDataAccess, demonstrate the use of the database class interface. We use the DynamicDataAccess example to demonstrate various techniques in the remainder of this chapter. You may have used the DynamicSQL sample application to verify your database connection. It is mentioned in [“Testing a Resource Name with the DynamicSQL Example” on page 48](#). You can find general instructions for both examples in [Appendix A, “Database Example Applications.”](#)

DBSession Methods

The DBSession class (along with some other classes) provides the methods you use to execute SQL in your TOOL code. In particular, you will invoke the following DBSession methods on the DBSession object for your database session:

Method	Description
CloseCursor	Closes a cursor opened by the OpenCursor or ExtendCursor method.
CloseExtent	When you are working with multiple result sets, this method cancels the current result set so you can use the ExtendCursor method to get to the next result set (Microsoft SQL Server and Sybase only).
Execute	Executes a prepared SQL DML statement (other than select or execute procedure).

Method	Description
ExecuteImmediate	Executes any single SQL statement that does not return results.
ExtendCursor	When you are working with multiple result sets, this method moves the cursor to the next result set or to the return parameters for the procedure (Microsoft SQL Server and Sybase only).
FetchCursor	Retrieves a single row or a set of rows from the result set of a cursor.
OpenCursor	Executes a prepared select or execute procedure statement to produce a result set.
Prepare	Prepares a SQL DML statement for execution, allowing you to use placeholders.
PreparePositioned	Prepares a SQL DML statement that references another cursor.
RemoveStatement	Removes a prepared statement.
Select	Executes a prepared select statement and retrieves a set of rows.

Executing Single SQL Statements

The simplest way to execute SQL is to use the `ExecuteImmediate` method. The `ExecuteImmediate` method executes a single SQL statement—of any type—stored in a string or a `TextData` object. The only restrictions are that the statement must not contain any placeholders or return any results. (Your particular database vendor may impose additional restrictions.) The `ExecuteImmediate` method is useful when you need to execute a statement only once.

```
DefaultDBSession.ExecuteImmediate('create table Paintings(
    Name varchar(60), Artist varchar(60), YearPainted integer);
```

The `ExecuteImmediate` method provides the same functionality as the `sql execute immediate` TOOL statement. There is no particular advantage to using `ExecuteImmediate` instead of the corresponding TOOL statement.

Using Prepared Statements

A prepared SQL statement is one that has been parsed by the database compiler. Prepared statements offer several advantages:

- You can execute the prepared statement multiple times (without reparsing), which improves performance for the second and subsequent iterations.
- You can use placeholders in the SQL statement.
- You can use a cursor to work with rows returned by the executed statement.

When you work with SQL statements that require input values (for placeholders) or that return results, you will use objects of the `DBDataSet` class in the `GenericDBMS` library. You will use `DBDataSet` objects both to describe and to store data for input values and result values. Some of the following examples demonstrate the use of `DBDataSet` objects.

There is a distinction between how you execute prepared `select` statements and prepared `insert`, `update`, or `delete` statements. While the initial steps for processing both types of statements are the same, `select` statements are executed using the `OpenCursor` or `Select` methods, and the other statements are executed using the `Execute` method.

Executing Prepared Queries

When you use the database class interface to execute a SQL statement that returns results, you may open a cursor to process the results a row at a time, or you may retrieve the entire result set all at once using the `Select` method.

The following steps summarize how to use a cursor to process SQL `select` statements. Each step is explained in detail below using examples from the `DynamicDataAccess` sample application.

► To execute a prepared `select` statement

1. Build the SQL statement.
2. Prepare the statement.
3. Set up a `DBDataSet` object with input values for placeholders, if necessary.
4. Open the cursor.
5. Fetch the rows.

6. Transfer or process the data.
7. Close the cursor.

About the DynamicDataAccess Example

The DynamicDataAccess sample application provides a window that lets the end user query the database by selecting specific criteria. The application constructs the actual SQL queries and then executes them using the appropriate GenericDBMS classes.

The DynamicDataAccess example allows end users to select or insert data from a database. [Figure 5-1](#) shows the Select screen for this example, as an end user might fill it in.

Figure 5-1 Choosing Selection Criteria

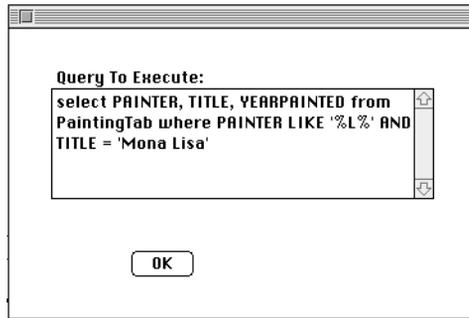
The screenshot shows a window titled "DYNAMIC DATA ACCESS". Inside, there are several controls:

- A "Table" dropdown menu set to "PaintingTab".
- A "Columns To View" list box containing "PRINTER", "TITLE", and "YEARPAINTED".
- A "Criteria List" table with columns: Column, Operator, Value, and AndOr.
- Buttons for "Execute", "Show Query", and "Quit" at the bottom.

Column	Operator	Value	AndOr
PRINTER	LIKE	%L%	<input checked="" type="radio"/> AND <input type="radio"/> OR
TITLE	=	Mona Lisa	<input checked="" type="radio"/> AND <input type="radio"/> OR
YEARPAINTED	=		<input checked="" type="radio"/> AND <input type="radio"/> OR

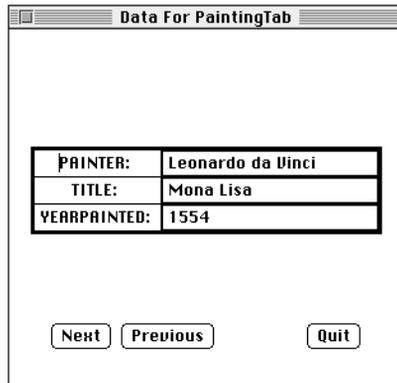
First the user chooses the table to query. Based on the table, the appropriate columns populate the Columns to View drop list and Criteria List array. Then the user chooses the columns to view and the criteria to use to determine the result set. When the user clicks the Execute button, iPlanet UDS constructs the `sql select` statement based on the criteria just entered. In this case, iPlanet UDS constructs the following SQL statement:

Figure 5-2 Dynamically Constructed SQL Statement



When the query executes, the application builds a screen that displays the data, as shown in [Figure 5-3](#):

Figure 5-3 Displaying the Results



The following section discusses the SQL methods used in this application.

Building the SQL Statement

When you write an application that will generate SQL statements at runtime, you must determine the range of information that the end user may require, write a user interface that allows the user to specify the full breadth of this information, construct valid SQL statements in your methods, perform the desired processing, and possibly display or return results. To construct a SQL statement, you simply build each phrase of the query as you collect input from the user.

The following method from the DynamicDataAccess application constructs a select statement based on input from the end user, as shown in [Figure 5-1](#). First, the BuildQuery method builds the select list:

```

CommandString = 'select ';

Selected : Array of ListElement;
Selected = <ColumnList>.GetElementList(ET_SELECTED);

NumColumns = Selected.Items;
for i in 1 to Selected.Items do
  if i < Selected.Items then
    Commandstring.concat(Selected[i].TextValue).concat(', ');
  else
    Commandstring.concat(Selected[i].TextValue);
  end if;
end for;

Commandstring.concat(' from ').concat(TableList);

Project: QueryMgr • Class: QueryWindow • Method: BuildQuery

```

Then the BuildQuery method determines if there is a where clause:

```

HasWhere : Boolean = FALSE;
Columns : integer = 0;
WhereClause : TextData = new;
WhereClause.concat(' where ');
for i in 1 to CriteriaArray.Items do

  if CriteriaArray[i].Value.IsNotEqual('') then
    WhereClause.concat(CriteriaArray[i].Column).concat(
      ' ').concat(CriteriaArray[i].Operator).concat(' ');
    WhereClause.concat('').concat(CriteriaArray[i].
      Value).concat('');
    if i < LastRow then
      WhereClause.concat('
        ').concat(CriteriaArray[i].AndOr).concat(' ');
    end if;
    HasWhere = TRUE;
  end if;
end for;

Project: QueryMgr • Class: QueryWindow • Method: BuildQuery

```

Finally, the `BuildQuery` method concatenates the `where` clause to the `select` statement:

```
if HasWhere then
    CommandString.concat(WhereClause);
end if;
```

Project: QueryMgr • **Class:** QueryWindow • **Method:** BuildQuery

Preparing the Statement

Before you can invoke the `OpenCursor` method, you must use the `Prepare` method to prepare the SQL statement. The `Prepare` method parses the statement and gives it a unique identifier so you can later reference it in the `OpenCursor` method. The `Prepare` method also returns an integer that indicates the type of statement being prepared. If you are using placeholders, the `Prepare` method returns information about the placeholders used in the statement. The following example from the `DynamicDataAccess` application illustrates preparing a SQL statement:

```
begin
dynStatement          : DBStatementHandle;-- Statement from Prepare
inputDescriptor       : DBDataSet;      -- Description of input
outputDescriptor      : DBDataSet;      -- Description of output
outputData            : DBDataSet;      -- Actual output
statementType         : integer;
rowType               : integer;
DataFetched           : StoredData = new;

    dynStatement = MySession.Prepare(
        commandString = CommandString,    -- the command
        inputDataSet = inputDescriptor,  -- output of substitutions
        cmdType = statementType);        -- output of statement type
```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** RetrieveData

This code returns an object of type `DBStatementHandle`, which is used to uniquely identify the prepared statement in subsequent method invocations.

Opening the Cursor

To execute the prepared `select` or `execute procedure` statement, use the `OpenCursor` method. This method opens the cursor and positions it before the first row in the result set. The `inputDataSet` parameter is required in the `OpenCursor` method, but it is not actually used in this example because there are no placeholders. A description of the expected output is given in the `resultDataSet` parameter. Notice that this method does not actually bring any data back in the `resultDataSet`, but only describes the data that will be returned with the `FetchCursor` method.

```
rowType = MySession.OpenCursor(
    statementHandle = dynStatement,      -- the statement
    inputDataSet = inputDescriptor;
    resultDataSet = outputDescriptor);  -- the output data
```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** RetrieveData

Before you retrieve the rows, the application must determine what columns will be retrieved so it can display them properly, as shown in [Figure 5-3](#). The example below uses the `GetColumn` method on the `OutputDescriptor` (which contains a description of the data) to retrieve the column description into `ColumnInfo`. It also stores a copy of the name in `ColumnsArray`, which creates the fields for the dynamically created window.

`ColumnInfo` is a `DBCColumnDesc` object. The `DBCColumnDesc` class describes a column in a relational table (or a single placeholder, as shown later in this section).

```
for j in 1 to NumColumns do
    ColumnName : TextData;
    ColumnInfo : DBCColumnDesc;
    ColumnInfo = OutputDescriptor.GetColumn(position = j);
    ColumnName = ColumnInfo.Name.Clone(TRUE);
    DataFetched.ColumnsArray.AppendRow(ColumnName);
end for;
```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** RetrieveData

Fetching Rows from the Result Set

To retrieve data from the result set, invoke the `FetchCursor` method. (An example appears in the next section.) `FetchCursor` retrieves one row or an array of rows from the result set and stores the values in the `resultDataSet` output parameter (a `DBDataSet` object). You can then work with these values in your code using the `GetValue` method of the `DBDataSet` class. Each time you invoke the `FetchCursor` method, the cursor moves forward either one row or the number of rows you specify.

Storing the Data

The `FetchCursor` method uses a `DBDataSet` object to store the fetched row or rows. Because iPlanet UDS uses the same `DBDataSet` object each time the `FetchCursor` method is invoked, if you need to preserve the data you should assign the `DBDataSet` values to other objects before the next `FetchCursor`, as in this example. For each row in the result set, the `GetValue` method retrieves the column values and appends them to the `DataArray` array. Remember, the column names have already been stored—the column values are matched by position.

```

numrows : integer;
totalRows : integer = 0;
while TRUE do
  numrows = MySession.FetchCursor(
    statementHandle = dynStatement,
    resultDataSet = outputData);
  if numrows <= 0 then
    exit;
  end if;
  TotalRows = TotalRows + 1;
  DataFetched.ResultSetArray[TotalRows] = new;
  for j in 1 to NumColumns do
    ColumnValue: TextData = new;
    (outputData.GetValue(position = j, value=columnValue));
    DataFetched.ResultSetArray[TotalRows].dataArray.
    AppendRow(ColumnValue);
  end for;
end while;

```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** RetrieveData

Closing the Cursor

After you have finished using the cursor, use the `CloseCursor` method to close it.

```
mySession.CloseCursor(statementHandle = dynStatement);
```

Executing Prepared DML Statements

You can use the `DBSession` class interface to execute `insert`, `update`, and `delete` SQL statements. This is particularly useful when you know what kind of statement you want to execute repeatedly, but you need to vary the criteria. iPlanet UDS methods use placeholders, which are mechanisms through which you can pass data into your methods at runtime—placeholders are discussed in detail later in this section.

The following steps summarize the use of the `DBSession` class methods to perform `insert`, `update`, and `delete` statements. Each step is described in detail below.

- **To execute a prepared insert, update, or delete statement**
 1. Build the SQL statement.
 2. Prepare the statement.
 3. Set the input values for placeholders in a `DBDataSet` object, if necessary.
 4. Execute the statement.
 5. Remove the statement if it will not execute again.

About the `DynamicDataAccess` Example

The `DynamicDataAccess` application includes a form through which users can insert data into the database, as shown in [Figure 5-4](#).

Figure 5-4 Insert Screen

Because this screen allows inserts only, the `sql insert` statement is already built in the example. This application uses placeholders to allow users to construct unique insert statements. These insert statements can also be built at runtime—see [“Building the SQL Statement” on page 122](#). Similarly, placeholders can be used in SQL statements built at runtime.

The application first checks to see if the user selected `PaintingTab` or `ArtistTab`. Based on the user’s choice, the appropriate fields are displayed. The user then enters the data to insert; this data is handled by the placeholders, described below.

Building the SQL Statement

In this particular example, we know we are building a `sql insert` statement. In the following SQL insert statement, notice that colons precede the placeholder names.

```

CommandString.setValue
('sql insert into PaintingTab (Painter, Title, YearPainted)
  Values (:Painter, :Title, :YearPainted)')

```

Project:DynamicDataAccess • **Class:** InsertDataWindow • **Method:** BuildInsert

Once the statement is built, it is ready to be prepared.

Preparing the Statement

Before you can invoke the `Execute` method, you must use the `Prepare` method to prepare the SQL statement. The `Prepare` method parses the statement and gives it a unique identifier so you can later reference it in the `Execute` method. The `Prepare` method also returns an integer that indicates the type of statement being prepared. If you are using placeholders, the `Prepare` method returns information about the placeholders used in the statement—this example does use placeholders, which are discussed below.

The following code demonstrates preparing a statement:

```
inputDescriptor : DBDataSet;
inputDescriptor = new;
dynStatement : DBStatementHandle;
statementType : Integer;
dynStatement = MySession.Prepare(
    commandString = commandString,
    inputDataSet = inputDescriptor,
    cmdType = statementType);
```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** RetrieveData

The `commandString` parameter contains the statement that was used in the previous step. The `inputDataSet` output parameter contains information about the placeholders (described below), and the `cmdType` parameter returns the type of statement.

Value	Description
DB_CV_INSERT	Insert statement
DB_CV_UPDATE	Update statement
DB_CV_DELETE	Delete statement
DB_CV_SELECT	Select statement
DB_CV_EXECUTE	Execute procedure statement

Finally, the `Prepare` method returns an object of type `DBStatementHandle`, which is used to uniquely identify the prepared statement in subsequent method invocations.

Processing Placeholders

A *placeholder* is an arbitrary name used in a SQL statement to represent a value that will be supplied when the statement is actually executed. The placeholder values are supplied through the `inputDataSet` parameter (described below) to the `Execute` or `OpenCursor` methods. For example, in the `DynamicDataAccess` example, the user can supply the data to insert multiple rows without the programmer having to re-prepare each `sql insert` statement.

NOTE In ANSI SQL, placeholders are represented by the “?” character. In iPlanet UDS, placeholders are represented by a colon preceding the variable name, as in “:name”. A placeholder name can be any legal iPlanet UDS name. However, if the name is the same as a SQL key word, you must enclose the placeholder name in quotation marks.

DBDataSet class You use the `DBDataSet` class to represent placeholders. The `DBDataSet` class describes the columns and the data in the placeholders and their values.

When you use the `Prepare` method to prepare the statement, the `inputDataSet` output parameter contains information about the placeholders that were included in the DML statement. This information is provided in the form of a `DBDataSet` object. You can then use the `SetValue` method to set the values of the individual placeholders for use in a subsequent `Execute` method. You can use the `SetValue` method to set values by name or by position; setting the values by position is more efficient.

The following example demonstrates setting the values for the placeholders referenced by the `inputDataSet` parameter.

```
inputDescriptor.SetValue(columnName = ':Painter',
    value = paintingData.Painter);
inputDescriptor.SetValue(columnName = ':Title',
    value = paintingData.Title);
inputDescriptor.SetValue(columnName = ':YearPainted',
    value = paintingData.YearPainted);
```

Project: QueryMgr • **Class:** ExecuteSQL • **Method:** ExecuteInsert

The `columnName` parameter to the `SetValue` method is set to the placeholder name; the `paintingData` object contains the data entered by the user and used to fill in the placeholders.

Executing the Statement

After you set the values for the placeholders, you can execute the statement. Use the `inputDataSet` parameter of the `Execute` method to reference the `DBDataSet` object that contains the values for the placeholders. In the following example from the `DynamicDataAccess` application, the `inputDataSet` parameter references the `DBDataSet` object whose values were set in the previous example:

```
MySession.Execute(statementHandle = dynStatement,  
    inputDataSet = inputDescriptor);
```

Note that the `Execute` method returns the number of rows affected, if known.

Removing the Statement

Use the `RemoveStatement` method when you no longer intend to execute a statement and you are finished with its result set. This is especially important if your application is complex and contains a number of statements, since each prepared statement consumes iPlanet UDS resources and usually database resources as well. It may also be necessary to use exception handlers to remove statements when exceptions occur.

```
MySession.RemoveStatement(statementHandle = dynStatement);
```

To see the `RemoveStatement` method in context, see the `Isql.Process` command in the example program `DynamicSQL` in [“DynamicSQL” on page 164](#).

Improving Application Performance

When you build an iPlanet UDS application that accesses one or more databases, you have a number of variables to consider when tuning performance. Applications can rapidly become very complex, as they scale up to accommodate new users, new components, and more data. While tuning an iPlanet UDS database application is an appropriate topic for a large book, this section contains some high level reminders that, in general, should always be heeded to maximize performance.

➤ **To maximize your application's performance, remember to**

1. Use explicit iPlanet UDS transactions.

For best performance you should use explicit iPlanet UDS transactions for all SQL statements, including selects.

2. Remove a cursor after closing it.

Unremoved cursors slowly use up the open cursors allowed by a database, and will eventually lock up the instance.

3. Wrap exception handlers around all of the prepare/fetch/close code and perform the remove in the exception handler as well. Otherwise exceptions occurring after the prepare will leave cursors unremoved.

Multi-threaded Database Access

iPlanet UDS supports non-blocking database connections on platforms where iPlanet UDS uses *native* threads and the vendor's libraries are thread safe. This multi-threaded database access is available on the following platforms:

- Windows platforms with a server partition connecting to a Microsoft SQL Server or Oracle database
- Sun Solaris, RS6000 AIX, and HP9000 HPUX platforms with a server partition connecting to an Oracle database

Additional platforms and databases vendors may also be supported for multi-threaded database access. Contact technical support for more information.

NOTE Multi-threaded database access does not apply to iPlanet UDS-threaded partitions. Partitions implementing iPlanet UDS threads must use serial access to databases.

For UNIX platforms only, any partition that participates in multi-threaded database access must be set to use native threads. This partition configuration parameter can be set as a partition property in the Partition Workshop, in Econsole for the installed partition, or using the Escript `SetThreadPkg` command.

When multi-threaded database access is enabled, an iPlanet UDS server partition can have other tasks running while one task is waiting on results from the database. Thus, accessing the database from one task does not block other tasks in the partition from running. This capability may allow sites to reduce the number of replicates needed when load balancing, or eliminate the need for process-level load balancing.

Note, however, that this does not change the transactional nature of `DBSession` objects. A `DBSession` object can only participate in one transaction at a time. Thus while the partition may not be blocked and may be able to run other tasks, those tasks may not be runnable if they reference a locked resource—for example, a `DBSession` object being used by another task in another transaction.

Enabling Single-Threaded Access to Databases

By default, multi-threaded database access is enabled for native-threaded partitions. To implement serial database access on UNIX platforms with native-threaded partitions, set the `FORTE_CLIB_LOCK` environment variable to `TRUE` (disabling multi-threaded database access).

Mapping `DBDataSet`s into Objects

When you select data, it is initially returned in a `DBDataSet` object. While it is most efficient to use the returned object directly, it is often not very convenient. Usually you will put the selected data into one or more custom objects. You typically use the `GetValue` method of the `DBDataSet` class to move data from the `DBDataSet` to your object.

While the `GetValue` method can handle any data type, it has four general variations. Performance for these variations varies. The variations are described below:

- `GetValue` can take either an integer *position* or *name* of the column for which to retrieve data.

Using positions is faster since iPlanet UDS can index the column directly rather than looking up a name in a list.

- GetValue can return a DBDataSet *return value* or pass an *output parameter*.

Using passed parameters is faster when you have an object into which you want to retrieve the result. Using a return value requires that a new object be created to hold the value; then the new object is returned. If you will then store the value in yet another object (and discard the returned object), it is more efficient to use the variation of GetValue that returns the value in a parameter.

For example this is faster:

```
myObj : myObjType;  
// All examples assume that myObj gets initialized somehow  
dataset.GetValue(position = 1, value = myObj.myAttr);
```

than this:

```
myObj : myObjType;  
myObj.myAttr.SetValue(dataset.GetValue(position = 1));
```

which is faster than:

```
myObj : myObjType;  
myObj.myAttr.SetValue(dataset.GetValue(columnName = 'mycol'));
```

Of course this assumes that myAttr has already been instantiated. If it hasn't the following may be more convenient:

```
myObj : myObjType;  
myObj.myAttr = dataset.GetValue(position = 1);
```

Note: the preceding is a little different than the seemingly equivalent:

```
myObj : myObjType;  
myObj.myAttr = new;  
dataset.GetValue(position = 1, value = myObj.myAttr);
```

The difference lies in the fact that the datatype of myAttr and the column in the dataset may differ. In the latter example the GetValue method will convert the value in the dataset to the type of myAttr (if possible). In the former example you will get a type clash error at runtime if the type of myAttr is not the same class or a superclass of the type of the column.

Vendor-Specific Information

Informix

Scroll Cursor Support

A scroll cursor must be used in an explicit iPlanet UDS transaction; it cannot be used in an implicit transaction. To use a scroll cursor, use the optional parameters for the DBSession methods Prepare and FetchCursor.

An example of using Informix scroll cursor feature follows:

```
SelectStatement : DBStatementHandle = NIL;
InputDescriptor : DBDataSet;
OutputDescriptor : DBDataSet;
CmdType        : Integer;
MaxRows        : 3;
Rows           : Integer;

-- Prepare the statement by indicating that we want
-- to open a scroll cursor.
SelectStatement = DefaultDBSession.Prepare(
    commandstring = 'select * from test',
    inputDataSet = InputDescriptor,
    cmdtype = CmdType,
    forScroll = TRUE);

begin transaction
--Scroll cursor can only be used within an explicit transaction.
DefaultDBSession.OpenCursor(
    statementHandle = SelectStatement,
    inputDataSet = InputDescriptor,
    resultDataSet = OutputDescriptor);

-- skip 2 records, fetch the next
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor,
    maxrows = MaxRows,
    key = DB_FK_RELATIVE, offset=2);
printResult(OutputDescriptor);
-- The method printResult is not showed here in the example.

-- fetch the next row
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor,
    maxrows = MaxRows,
    key = DB_FK_CURRENT);
```

```
printResult(OutputDescriptor);

-- fetch the 8th row
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor, maxrows = MaxRows,
    key = DB_FK_ABSOLUTE, offset=8);
printResult(OutputDescriptor);

-- fetch the previous (7th) row
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor, maxrows = MaxRows,
    key = DB_FK_PREVIOUS);
printResult(OutputDescriptor);

-- fetch the first row
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor,
    maxrows = MaxRows,
    key = DB_FK_FIRST);
printResult(OutputDescriptor);

-- fetch the next (2nd) row
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor,
    maxrows = MaxRows,
    key = DB_FK_NEXT);
printResult(OutputDescriptor);

-- fetch the last row
Rows = DefaultDBSession.FetchCursor(
    statementHandle = SelectStatement,
    resultDataSet = OutputDescriptor,
    maxrows = MaxRows,
    key = DB_FK_LAST);
printResult(OutputDescriptor);

DefaultDBSession.CloseCursor(statementHandle = SelectStatement);
end transaction;
```


Transactions

When you include SQL statements (or equivalent methods) in an iPlanet UDS application, you use the iPlanet UDS transaction model to start explicit transactions or under some circumstances, implicit transactions.

This chapter describes how to use transactions so that your application uses transactions properly and with the best possible performance. A rule of thumb is to always use explicit iPlanet UDS transactions for best performance.

This chapter discusses a variety of transaction scenarios, and includes the following topics:

- explicit and implicit transactions
- dependent and independent transactions
- common problems with shared and transactional objects
- two-phase commit and distributed transactions
- vendor-specific notes

Relationship Between iPlanet UDS and Database Transactions

When you write an iPlanet UDS application that accesses a database, you will structure your application into iPlanet UDS *transactions*. (iPlanet UDS transactions are the same as TOOL transactions.) Like a database transaction, an iPlanet UDS transaction should consist of a set of statements that should be treated as a unit, succeeding or failing together.

The only transaction commands you use are to start and end iPlanet UDS transactions. You do *not* use the explicit transaction statements for your database (such as `commit`, `rollback`, or `set autocommit`), because iPlanet UDS starts and coordinates all database transactions on your behalf. You can, however, use the iPlanet UDS Abort method to abort both the iPlanet UDS and DBMS transactions. Thus, you need only be concerned with iPlanet UDS transactions, because all database transactions are transparently managed.

You should always use explicit transactions. If you do not, iPlanet UDS will start implicit transactions and your application's performance may be significantly impacted. The following sections describe explicit and implicit transactions in more detail.

A transaction that is entirely local provides better performance. A transaction that starts and ends in the same partition and makes no intervening remote calls is a local transaction. A transaction becomes distributed when a remote call is made. To the degree that you can increase the number of local transactions in an application relative to distributed transactions, you can improve the performance of the application.

Explicit iPlanet UDS Transactions

You should always include SQL statements and methods in explicit iPlanet UDS transactions. Your application will perform far better if you do so, particularly when you select data using a cursor (this includes `TOOL sql select` statements where the target is an array object). In addition, some operations behave differently when not used in an explicit transaction (see the iPlanet UDS online Help).

Starting an iPlanet UDS transaction To start an explicit iPlanet UDS transaction, you can do either of the following:

- use the `begin transaction` `TOOL` statement (described in the *TOOL Reference Guide*)
- use the `TransactionHandle` class (described in the iPlanet UDS online Help)

Ending an iPlanet UDS transaction To end an iPlanet UDS transaction, you use the `end transaction` clause of the `begin transaction` statement, or you can use methods on the `TransactionHandle` class. An iPlanet UDS transaction is always ended explicitly; no type of event or statement will implicitly end an iPlanet UDS transaction.

Any number of SQL statements can be contained in a transaction; however, the transaction should consist only of statements that should succeed or fail together — no other statements. You should also exclude any SQL statements that your database does not allow within a transaction (for example, some systems do not allow DDL statements in transactions).

The following example shows two SQL statements enclosed in an explicit transaction defined using TOOL:

```
begin transaction
  sql insert into mytab (a) values (10) on session MySession;
  sql insert into othertab (b) values (12) on session
YourSession;
exception
  when e : GenericException do
    task.ErrorMgr.showErrors(TRUE);
end transaction;
```

When you embed SQL in an explicit iPlanet UDS transaction, iPlanet UDS automatically starts a database transaction, and both the iPlanet UDS and database transaction must succeed together or both will fail. That is, if failure occurs in the underlying database transaction, it rolls back and the enclosing iPlanet UDS transaction also aborts. Or, if the enclosing iPlanet UDS transaction aborts, iPlanet UDS rolls back the underlying database transaction.

Implicit iPlanet UDS Transactions

If you invoke a method on the DBSession object or issue a TOOL SQL statement that accesses the database, and no iPlanet UDS transaction is currently in effect, then iPlanet UDS starts an implicit iPlanet UDS transaction. While an implicit transaction does ensure the integrity of your data with respect to the database, it can also have a severe impact on performance. An implicit transaction is slower and may consume more memory than an explicit transaction.

An implicit iPlanet UDS transaction always has a scope of one SQL statement. Execution of multiple SQL statements enclosed in separate transactions is generally slower than if the statements are executed within a single transaction. While this may be acceptable for testing or ad hoc use, implicit transactions should never be used in deployed applications because they dramatically reduce performance.

Single SQL DML Statements

If you issue a single SQL statement, such as `sql delete`, `sql insert`, `sql update`, and `sql execute immediate`, iPlanet UDS places the statement in an implicit iPlanet UDS transaction encompassing just the SQL statement. That is, the transaction will be started before the SQL statement and committed after it. No other TOOL statements are included in the transaction. For example:

```
-- An insert
sql insert in paintings values (:newPainting);
```

SQL Execute Immediate Statements with DDL

The `sql execute immediate` statement can be used to issue data definition language (DDL) statements, such as `create table`. The database vendors treat DDL statements differently with respect to transactions, as shown in the following table:

Database	Interaction between a DDL statement and a transaction
DB2	Allows one or more DDL statements in a multi-statement transaction.
Informix	Allows one or more DDL statements in a multi-statement transaction.
ODBC	Uses the semantics of the underlying database.
Oracle	Use of a DDL statement automatically commits the current transaction.
Rdb	Allows one or more DDL statements in a multi-statement transaction.
Sybase	Allows only one DDL statement in a transaction.

If you need to ensure portability across databases, do not put `sql execute immediate` statements containing DDL in iPlanet UDS transactions that contain more than one SQL statement. Instead, enclose each DDL statement in an independent transaction block.

Use of Cursors in Implicit iPlanet UDS Transactions

While you can use cursor statements in implicit transactions, it is not recommended. iPlanet UDS buffers the entire result set for a cursor, which consumes memory. As a result, performance may degrade because iPlanet UDS retrieves the entire result set before executing the next statement. The following sections describe the effect of implicit iPlanet UDS transactions on the TOOL SQL statements `sql open cursor`, `sql fetch cursor`, and `sql close cursor`. In general, the descriptions below also apply when you use the `OpenCursor`, `FetchCursor`, and `CloseCursor` methods on a `DBSession` object.

SQL Open Cursor

A `sql open cursor` statement executed outside an explicit TOOL transaction is enclosed in an implicit transaction. If a cursor is opened in an implicit transaction, the `sql open cursor` statement alone causes all the data to be fetched, buffered, the database cursor to be closed, and the implicit transaction to end. The database cursor is closed because most databases do not allow a database cursor to remain open after the transaction ends. However, the TOOL cursor remains open and subsequent `sql fetch cursor` statements retrieve results from the cursor as usual, but the data is retrieved from the buffer, rather than directly from the database.

Executing `sql open cursor` within an implicit transaction has several effects, all of which can degrade performance:

- The entire result set is retrieved for the cursor, and data is copied twice: once into the buffer, and again into wherever you finally put it.
- It requires additional memory, as the entire result set must be buffered in memory at one time.
- Control is not returned back to the user to start processing the first row of data until all of the data is retrieved from the database.

In addition, because the implicit transaction automatically buffers all the data and commits the transaction, the actual data in the database corresponding to the buffered data may be concurrently updated without any notification to your TOOL program, or any locking on the data. This is not a problem for read-only data, but for updatable data, you should be sure to consider a concurrency strategy that is appropriate. For maximum consistency, use explicit iPlanet UDS transactions.

SQL Fetch Cursor

The `sql fetch cursor` statement is never enclosed in an implicit transaction. The behavior of a `sql fetch cursor` statement executed outside of an explicit iPlanet UDS transaction depends on when (how) the cursor was opened:

- If the cursor was opened in an explicit iPlanet UDS transaction, then the cursor was closed at the end of the transaction; any fetch outside the transaction in which the cursor was opened will fail.
- If the cursor was opened outside of an explicit transaction, the result set for the cursor was buffered and the fetch statement will fetch the next row from the buffer.

The following example shows two cursors: `cursor1` is opened in an implicit transaction and `cursor2` in an explicit transaction:

```
sql open cursor1...-- surrounded by an implicit transaction

begin transaction
sql open cursor2...
while condition do
  sql fetch cursor2...-- fetches a row from the database
  sql fetch cursor1...-- fetches a row from the buffer
  other TOOL statements;
end while;
end transaction;

sql fetch cursor2...
-- This will cause an exception because the
-- cursor was closed at the end of the transaction.

sql fetch cursor1...
-- This will execute successfully because it continues
-- to fetch from the buffer.
sql close cursor1...
-- This will execute successfully and change the state of
-- the cursor to closed so it may subsequently be opened again.
```

NOTE When using the `FetchCursor` method on the `DBSession` object on a cursor that was opened in an implicit transaction, you should not use the `maxrows` input parameter to determine the number of rows returned. Instead, check the functional result. Generally rows are returned one at a time.

For Loops

When you use a `for` loop with a cursor name or `sql select` statement, iPlanet UDS performs an open cursor as part of the initialization of the `for` loop. If no explicit iPlanet UDS transaction is in effect, this open cursor is enclosed in an implicit transaction, causing the result set to be buffered as described earlier in the “SQL Open Cursor” section. However, the body of the `for` loop is not part of the transaction.

SQL Close Cursor

When you issue the `sql close cursor` statement against a buffered cursor, the data buffered for the cursor is released.

Independent, Dependent, and Nested Transactions

It is strongly recommended that you always use transaction blocks for units of work that might be a transaction. When you start an explicit transaction, you can make it one of three types of transactions:

- independent
- dependent
- nested

For transactions that interact with databases, you will most often use dependent, and in some cases independent. (Dependent transactions are the default if you do not specify a type in the `begin transaction` statement.) You should not declare nested transactions that manipulate database data (see below for an explanation).

Using Dependent Transactions

A coding practice that makes transaction design more modular is to define as a dependent transaction any segment of code which might comprise a unique transaction in some context. This has the dual benefits of ensuring that all SQL always runs in an explicit transaction with no constraint on its re-use. In fact, the default type of transaction is dependent; the statement `begin transaction` is the same as `begin dependent transaction`.

An example will help illustrate. Assume a method containing the following code fragment:

```
begin dependent transaction;  
SQL UPDATE personTable set dept = :var where dept = :cond;
```

This code can run as a top level transaction or it can be invoked from a context where a transaction is already in progress:

- If it is invoked when a transaction is not in progress, it simply starts a new transaction.
- If it is invoked when another transaction is in progress, it would just become part of the transaction already in progress.

However, if instead you used `begin independent transaction`, this method could not be invoked when a transaction was already in progress, since an independent transaction can not be nested inside another transaction.

Using Independent Transactions

Since dependent transactions can be used whether or not a transaction is in progress, when are independent transactions preferred? While dependent transactions are appropriate for many situations, there are circumstances when you should start a new transaction, rather than pick up or continue from a previous step. In such a case you should explicitly use an independent transaction (using the `TOOL` statement `begin independent transaction`, for example).

Avoid Nested Transactions

While the preferred way to execute all SQL statements (using `TOOL` or equivalent methods) in an iPlanet UDS application is within an explicit iPlanet UDS transaction, you should avoid placing SQL statements within *nested* iPlanet UDS transactions. While iPlanet UDS does not flag nested iPlanet UDS transactions containing SQL as an error, using this construction can result in inconsistent data for the following reasons.

While iPlanet UDS nested transactions provide a useful mechanism for fine-tuning the scope of a rollback in applications with complex transactional structure, most database vendors do not support nested transactions. The result of this fact is that SQL statements in a nested iPlanet UDS transaction are considered *part of the outer iPlanet UDS transaction*, not the nested transaction. (All other transactional operations within the nested transaction succeed or fail together as the nested transaction.) This has the following implications:

- If the nested iPlanet UDS transaction aborts, but the outer iPlanet UDS transaction commits, the SQL work is committed.
- Only if the *outer* transaction aborts is the SQL work aborted or rolled back.
- If the nested transaction aborts, the SQL statements within the nested transaction are not rolled back, as long as the outer transaction commits. However, other iPlanet UDS transactional logging which may have occurred within the nested transaction *is* rolled back, potentially leaving an inconsistent data state.

This behavior is true regardless of whether the nested transaction is invoked locally or remotely relative to the outer transaction, and whether the transaction is within a bracketed transaction block or asynchronous nested task.

Common Problems with Shared and Transactional Objects

Because the DBSession object is both a shared and transactional object, you must take the same precautions that you would for other shared and transactional TOOL objects. The most common problems are described below—note that all these problems can occur in normal TOOL objects with the same properties.

In general, you should avoid starting transactions before any operation that might wait—for example, before an event loop where your application might wait for user input. If you begin a transaction in such a situation, you might unnecessarily cause data to be locked, preventing access by other users.

Unexpected Blocking Due to a Long-Running Query

While executing a query in a database session, the DBSession object holds the internal mutex lock on the session. While this mutex is held, no other query may be executed. Note that due to the nature of iPlanet UDS's database interface, the entire iPlanet UDS server process may be blocked while the query is running.

Unexpected Blocking Due to a Long-Running Transaction

If a task begins a transaction and then executes a TOOL SQL statement against a DBSession object, that session's transaction lock is now exclusively held by the task's transaction. The lock remains held until the transaction terminates (commits or aborts). No other tasks may execute any SQL statements against that same session, unless those tasks are in the same transaction or the owning transaction terminates. For example, assume the following code fragment uses a DefaultDBSession:

```
begin transaction
  event loop
    SQL select...into array object;
    when <detail>.click do
      start task self.SelectCurrentRow(currentRow)
        where transaction = dependent;
    when <update>.click do
      SQL...
    when <save>.click do
      exit;
    when <abort>.click do
      transaction.Abort(TRUE)
    end event;
end transaction;
```

Note that an event loop begins after the transaction has begun. Because an event loop can potentially last a long time, any data that is locked by a given SQL statement will remain locked until you exit the window.

Any child transaction is dependent on the parent task's transaction; thus, any queries executed within the child task are not blocked. However, queries executed by other tasks not participating in the same transaction (or from other clients) are blocked.

Avoiding Deadlocks

It is possible under some scenarios to encounter a database deadlock. These scenarios most often entail multiple database sessions accessing the same table data on the same server. For example:

- If an application has two or more database sessions accessing the same table on the same server (but different rows), a deadlock may occur on a database that uses page-level locking (Sybase Adaptive Server 11.5, for example).

NOTE Starting with Adaptive Server 11.9, Sybase provides row-level locking.

- If an application has two or more database sessions that are not only accessing the same table on the same server, but also the same rows, then a deadlock may occur on a database that uses row-level locking (Oracle, for example).

The fundamental basis for the deadlock is that the database cannot detect that separate iPlanet UDS sessions are part of the same transaction. These scenarios may result in deadlocks, whether the database sessions originate from one or more clients, from one or more iPlanet UDS transactions, or from one or more partitions (or from multiple applications).

You can usually write an application so as to avoid deadlocks. Depending on the requirements of the application, the following hints may be useful:

- Use a single database session to access data. Multiple requests within the same transaction and same database session do not block.
- When using multiple database sessions, try to ensure that each session accesses different tables.
- When using multiple database sessions and accessing the same data, an application should minimize contention at the database level by synchronizing the invocation of transactions. This can be accomplished by using an iPlanet UDS transaction lock to block and defer other transactions at the iPlanet UDS level.

- When possible, narrow the scope of each transaction to encompass the fewest SQL statements using a single database session. This minimizes the possibility of deadlock, since locks held are immediately released upon completion of the explicit transaction.

If a database deadlock occurs, the database will raise a `DBDeadlockException` to the iPlanet UDS application. Your application should be prepared to handle this potential exception.

Transactions and Database Sessions

A single iPlanet UDS transaction can access more than one database session, more than one database, or more than one database resource manager. iPlanet UDS acts as the transaction coordinator for all iPlanet UDS transactional objects, coordinating all commits and aborts.

For example, if a transaction accesses two resource managers and one resource manager aborts the transaction, iPlanet UDS notifies the other database that the transaction is aborted. Likewise, when the transaction commits, iPlanet UDS signals all resource managers to commit the transaction.

While iPlanet UDS coordinates transaction commit across all of the sessions using a two-phase commit protocol, the iPlanet UDS database interface cannot currently use DBMS two-phase commit mechanisms. The iPlanet UDS transaction manager will commit all database transactions first, and then the recoverable object transactions. This means that you should plan for failures of multi-session commits accordingly.

Multitasking in a Database Session

If you want to use the same database session in multiple concurrent tasks, all the tasks must be in the same transaction. You can add a task to a transaction in progress in one of two ways:

- Use the `dependent` option on the `start task` statement.
- Use the `Join` method on the `TransactionHandle` class.

See the *TOOL Reference Guide* for more information on transactions and multitasking

iPlanet UDS Support for Two-Phase Commit

iPlanet UDS support for two-phase commit (2PC) is limited by general database vendor implementation restrictions on multiple-vendor, two-phase commit. iPlanet UDS guarantees two-phase commit with respect to iPlanet UDS objects; however, iPlanet UDS only guarantees two-phase commit with respect to database data if a third-party transaction monitor, like Encina or Tuxedo, is used in conjunction with iPlanet UDS. (See [“Two-Phase Commit with Multiple Database Vendors” on page 153.](#))

What is two-phase commit? The purpose of *two-phase commit* is to allow a transaction to span multiple databases, regardless of vendor, with a guarantee that the transaction is committed only after all databases involved in the transaction guarantee that the commit will succeed. For example, the following “code” would use two-phase commit:

```
begin transaction
sql update ... (on dbms 1);
sql update ... (on dbms 2);
end transaction;
```

iPlanet UDS two-phase commit implementation The main purpose of two-phase commit in a distributed transaction manager is to enable recovery from a failure that occurs *during the actual transaction commit processing*. The iPlanet UDS transaction manager was built with this in mind, but only with respect to the “volatile” (or “in memory”) objects that iPlanet UDS manages. What this implies is that because iPlanet UDS stores objects in memory and not persistently on disk, the requirement of recovery for these objects is significantly reduced.

In the iPlanet UDS distributed two-phase commit model, tasks and messages carry along with them transaction identification and, during commit processing, every distributed participant is polled for its availability to commit the transaction. If an application saves persistent data to disk during a distributed iPlanet UDS transaction, that application should address the potential for failure during the transaction commit processing.

The iPlanet UDS prepare phase polls each site (confirming a communications link with each distributed participant) but no prepare request goes to the database. When all sites are ready to commit, *iPlanet UDS expects that the commit will complete successfully*. Then, at this point, called the “commit window,” three scenarios can occur:

- All participating servers commit successfully. The distributed transaction is committed, and all data is consistent.
- One participating server terminates (with data not yet committed) but no other participating server has committed its unit of work. In this scenario, iPlanet UDS will abort the transaction, and the data will not be inconsistent. The entire transaction is, in effect, rolled back.
- One participating server terminates (with data not yet committed) while a second participating server has already committed its unit of work. In this scenario, the outcome of the distributed transaction is inconsistent data.

Only the last scenario is cause for concern. Without two-phase commit, if one server fails during the process of committing (that is, while `TOOL` is executing the `end transaction` clause), there is no guarantee that the other server(s) will abort, as they might have already committed. Therefore, if a transaction is being committed and it accesses multiple database sessions, multiple databases, or multiple DBMS resource managers and there is a failure during the commit window, your data might be inconsistent. You should take this into consideration when planning your transactions.

Mission critical applications that require distributed transactions can anticipate the third “commit window” scenario in a number of ways:

- Utilize a TP monitor such as Encina (see below)
- Log distributed updates in an auxiliary database table (much like a distributed transaction monitor's transaction-state log). While this approach has been the traditional banking application solution prior to the commercial availability of products like Encina, it is somewhat complex. Nor is it generic enough so as not to have to change code every time a new table or database site is introduced into the data model.
- Rearrange the data model in order to eliminate the need for distributed transactions. This is usually only a temporary solution (with smaller numbers of active clients) and cannot be applied to complex legacy systems.

With the advent of the X/Open distributed transaction architecture (the XA Interface) more database vendors have found that by complying with the XA interface they can plug their database-specific implementation of transaction into a globally managed transaction, with commit and abort processing being conducted by a central coordinator. Of course, the overall transaction manager coordinating the global transaction must itself, persistently record the state of the different distributed branches participating in the transaction. A significant portion of the functionality provided by products such as Encina is to provide exactly this global transaction management.

Two-Phase Commit with One Database Vendor

You can use a true implementation of two-phase commit in an iPlanet UDS application if both of the following conditions are true:

- The distributed transaction is between two or more databases from the same vendor.
- The distributed transaction is initiated by the SQL statements issued within one database, which in turn calls another database to complete the transaction.

In this scenario, the two-phase commit protocol of the RDBMS itself manages the transaction across the multiple databases; iPlanet UDS could be considered “outside” of the actual distributed transaction.

Two-Phase Commit with Multiple Database Vendors

Two-phase commit with multiple databases, except for the limited case above, requires the use of Encina with iPlanet UDS (see *Integrating with External Systems*). To use Encina you must use an XA-compliant TP monitor (Encina) along with the two XA-compliant database systems.

Notes on Vendor-Specific Transaction Handling

Generally, you should avoid using explicit database transaction statements, as iPlanet UDS starts database transactions for you. And, as described in “[SQL Execute Immediate Statements with DDL](#)” on page 142, you should only execute SQL DDL statements individually, in explicit independent transactions. However, some transaction control statements are allowed.

NOTE iPlanet UDS does not support the use of autocommit. Do not use the statement `set autocommit on` within a `sql execute immediate` statement.

DB2

DB2's prepared SQL statements do not persist across transaction boundaries. iPlanet UDS will automatically re-prepare `TOOL` statements as required. To minimize this overhead, you should use multi-statement, explicit transactions.

Informix

The iPlanet UDS interface to Informix IDS Server supports connections to each type of Informix database:

- a non-ANSI database with logging
- a non-ANSI database without logging
- an ANSI database

While Informix handles transactions and logging differently in these databases, iPlanet UDS transparently accesses any database, and issues (or ignores) transactional SQL statements as appropriate for that database type.

NOTE `NOT WAIT` is the default for Informix's `SET LOCK MODE`; any lock conflicts encountered will return an error immediately.

Rdb

If you use multiple database sessions within an iPlanet UDS transaction, take care to avoid locking conflicts between the sessions, because Rdb does not allow for the sharing of locks within an Rdb process. This is a concern only if the sessions are running within the same iPlanet UDS server partition.

By default, all Rdb transactions initiated through the iPlanet UDS application interface are in “READ WRITE NOWAIT” mode.

To set an Rdb database session to a different transaction mode, you can issue a `sql execute immediate` statement in your application code to issue Rdb `DECLARE TRANSACTION` or `SET TRANSACTION` statements. Because these statements are valid only if there is no active Rdb transaction, you must issue these statements immediately after a `TOOL begin transaction` statement. For example:

```
begin transaction
  sql execute immediate 'set transaction read only';
  ...
end transaction;
```

Alternatively, a `DECLARE TRANSACTION` statement issued implicitly (outside of a `TOOL` transaction) will affect the mode of a subsequent Rdb transaction, either implicit and explicit. For example:

```
-- Read-only SQL operations
sql delete from emp;
```

The `sql delete` statement will fail because its implicit transaction is in `READ ONLY` mode.

Error Handling

iPlanet UDS converts all database errors into exceptions so that a single exception handler can process all errors that occur in a set of SQL statements. This allows you to separate error handling code from the body of a program, providing for more readable and maintainable code.

This chapter contains a brief overview about how iPlanet UDS handles database exceptions.

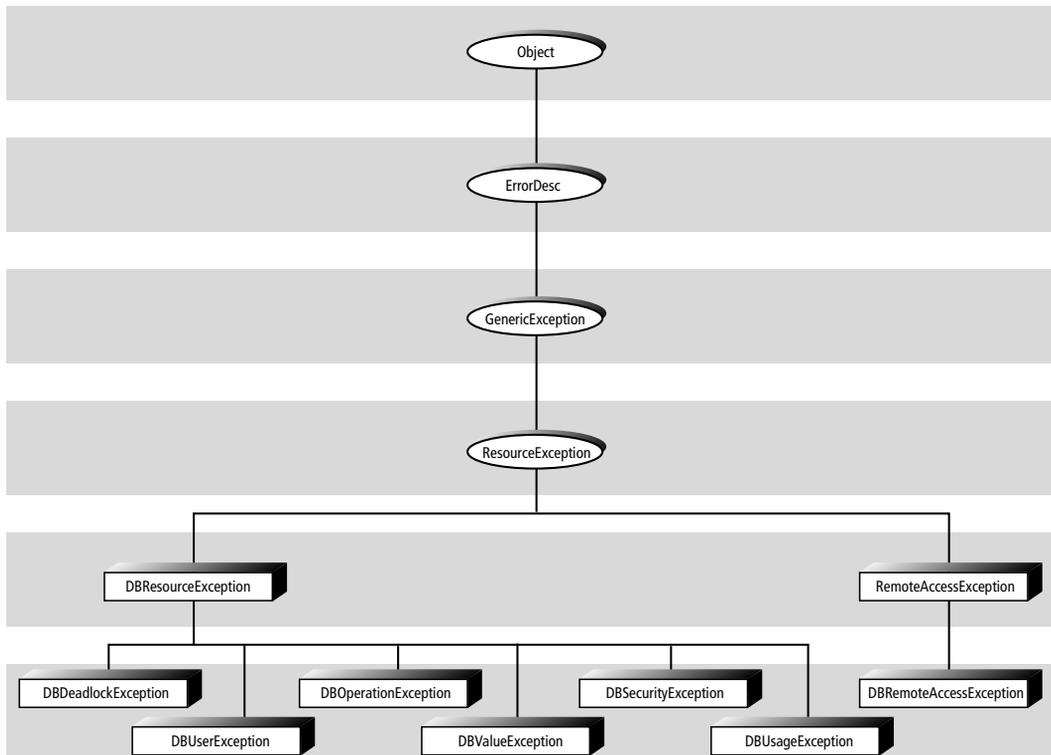
See the iPlanet UDS online help (under Native Database Errors) where you can find, for each supported database, a mapping of native database errors to iPlanet UDS exceptions. The online help also lists the ODBC and ISO SQL-92 mappings to iPlanet UDS exceptions.

The Exception class reference can be found in the iPlanet UDS online help.

Types of Database Exceptions

The class hierarchy for exceptions that may be returned in an iPlanet UDS database application is shown in [Figure 7-1](#).

Figure 7-1 The Database Exception Hierarchy



iPlanet UDS uses two general classes for database exceptions: `DBRemoteAccessException` and `DBResourceException`.

DBRemoteAccessException class Exceptions of class `DBRemoteAccessException` occur when a session cannot be initiated or a current session is no longer valid. These types of errors are generally non-recoverable in the sense that you lose the current transaction and session if one existed. You can recover in the sense that you can attempt to start a new database session, or connect to a different server. Refer to the iPlanet UDS online Help for more information on this exception class.

DBResourceException class Exceptions of class DBResourceException occur when accessing a database. These types of errors are considered recoverable, since the current transaction or session is not lost, and the application can continue after handling the exception. See the iPlanet UDS online Help for more information on this class.

Many common database errors result in exceptions of type DBResourceException, such as incorrectly formed SQL statements, data constraint violations, access violations, and so on. In fact, there are several subclasses of the DBResourceException class as shown below:

Exceptions	Cause
DBDeadlockException	A database deadlock.
DBOperationException	An error in the runtime operation of a SQL statement.
DBSecurityException	A security violation.
DBUsageException	Incorrect TOOL program use for the (particular) database: syntax errors, attempt to use non-existent tables, and so on.
DBUserException	An error generated by a user database procedure or trigger.
DBValueException	An error in assigning a value, either through conversion, truncation, constraint, or arithmetic error.

See the iPlanet UDS online help for more information on each of these subclasses.

The iPlanet UDS online help (Native Database Errors) contains tables that map vendor-specific errors to these exception classes. Refer to the table for your database vendor to become familiar with how errors generated in your database map to the various iPlanet UDS exceptions.

Database Example Applications

iPlanet UDS provides a number of example applications that illustrate how to use the features described in this manual. This appendix provides instructions on how to install the examples, a brief overview of the applications to help you locate relevant examples, and a section describing each example in detail. Typically, you run an example application, then examine it in the various iPlanet UDS Workshops to see how it is implemented. You can modify the examples if you wish.

How to Install iPlanet UDS Example Applications

You can access the iPlanet UDS example applications only if they have been installed into your central repository or into a private local repository during installation of iPlanet UDS, or if you have imported them into your repository.

The examples are located in subdirectories under the `FORTE_ROOT/install/examples` directory. The example applications are stored as `.pex` files. If they are not already installed in your repository, import them by including the `tstapps.fsc` script in `Fscript`. The `tstapps.fsc` script is located in the `FORTE_ROOT/install/examples/install` directory. Bring up `Fscript` in standalone mode and issue the following commands:

```
fscript> UsePortable
fscript> SetPath %{\FORTE_ROOT}/install/examples/install
fscript> Include tstapps.fsc
```

This will import most of the example applications and quit `Fscript`. Note that certain highly specialized examples are not automatically imported by `tstapps.fsc`.

To run an application, select it in the Repository Workshop's plan browser and then click on the Run button.

If you want to remove all the examples from your workspace, you can do so by including the `remprj.fsc` script in Fscript. Bring up Fscript in standalone mode and issue the following commands:

```
fscript> UsePortable
fscript> SetPath %{\FORTE_ROOT}/install/examples/install
fscript> Include remprj.fsc
```

This will exclude all the example applications and quit Fscript.

Overview of Database Example Applications

This section provides an overview of the database example applications. For a complete list of the iPlanet UDS example applications, see *A Guide to the iPlanet UDS Workshops*.

You can find the `.pex` files for the examples in the directory `FORTE_ROOT/install/examples/database`. For the complete description of an individual application, see "[Application Descriptions](#)" on page 163, which lists the applications in alphabetical order.

GenericDBMS Library Examples

Example	Description
DynamicDataAccess	Shows how to build dynamic SQL queries.
DynamicSQL	Illustrates the use of the GenericDBMS classes with a command line utility.
QueryMgr	Acts as a server for DynamicDataAccess.
WinDB	Illustrates the use of GenericDBMS classes in a window-based application.

Application Descriptions

This section lists the example applications in alphabetical order. Each example has five sections describing it.

The *Description* section defines the purpose of the example, what problem it solves, and what TOOL features and iPlanet UDS classes it illustrates.

The *Pex Files* section gives you the subdirectory and file names of the exported projects. The examples are in subdirectories under the FORTE_ROOT/install/examples directory. You can import example applications individually if you wish. When multiple .pex files are listed, there are supplier projects in addition to the main project. You will need to import all the files listed to run the application. Import them in the order given so that dependencies will be satisfied.

The *Mode* section indicates whether the application can be run in either standalone or distributed mode, or whether it must be run in distributed mode.

The *Special Requirements* section identifies whether you need a database connection, an external file, or any other special setup.

Finally, the *To Use* section tells you how to step through the application's functions.

See the *iPlanet UDS System Management Guide* if you need directions to set up an iPlanet UDS server. See [Chapter 3, "Making a Database Connection"](#) if you need information on how to make a connection to a database. The database examples run against either Sybase or Oracle.

DynamicDataAccess

Description DynamicDataAccess lets you construct ad hoc SQL queries. The data fields used to display the data and the text graphics used for labels are created at runtime to match the columns in the dynamic query.

Pex Files database/querymgr.pex, database/dda.pex.

Mode Distributed only.

Special Requirements Database connection. The files artist.dat and painting.dat must be located in FORTE_ROOT/install/examples/database.

➤ **To use DynamicDataAccess**

1. Before running DynamicDataAccess, open the Project Workshop for the QueryManager project.
2. Open the DBResourceMgr service object properties dialog. Make sure the Database Manager value is correct for your database connection.
3. Run DynamicDataAccess. You will be prompted for a User Name, Password, and Database. You must provide all three values.

This application will create its own tables and data in the database you selected. The data is read in from the files FORTE_ROOT/install/examples/database/artist.dat and painting.dat.

4. Choose Select from the radio list to select from existing artist and painting values. Choose Insert to add your own data. Follow the prompts in the Insert and Select Windows to construct SQL statements and view the results.

DynamicSQL

Description DynamicSQL is a command line utility illustrating the use of the GenericDBMS Library classes. It lets you perform standard SQL database access commands.

Pex Files frame/utility.pex, database/dynsql.pex.

Mode Distributed only.

Special Requirements Database connection, table creation (if necessary).

➤ **To use DynamicSQL**

1. You need an environment that has a node with a resource manager. Before running DynamicSQL, open the DynamicSQL Project Workshop. Open the AnyDBMgr service object property sheet. Make sure the Database Manager value is correct for your database connection.

- DynamicSQL does not create any tables. You must either create them ahead of time, or use an existing test table. If you need to create a test table, the files `maketst.syb` and `maketst.ora` are provided in `FORTE_ROOT/install/examples/database`. If you will be using `maketst.syb`, edit the first line to use an existing database. For example, create a database called `testapps`, then edit `maketst.syb` to start with:

```
use testapps
```

Use the standard mechanism for redirecting the `maketst` file to load the data into your database.

- Run DynamicSQL. You will need to know the name of the database you will use, and a `userid` and password. Make sure that the iPlanet UDS Launch Server - Console window is visible. It is this application's only window.
- At the prompts enter your database name, then your `userid` and password. You should then see a `SQL>` prompt. If you used the `maketst` file, enter:

```
SQL> select * from alltypes;
```

- You should see the data displayed. Then enter:

```
SQL> select * from alltypes where intcol = :a;
Enter values for the following placeholders:
a:> 1
Reexecute? (Y/N)y
Enter values for the following placeholders:
a:> 5
Reexecute? (Y/N)y
```

- To end your session, type:


```
SQL> exit
```
- The program will print out "Disconnecting," but you will need to return the focus to the iPlanet UDS Workshop yourself.

WinDB

Description WinDB uses the GenericDBMS Library classes. It lets you perform standard SQL database access commands. It also illustrates how to send and retrieve Binary Data (BLOBs) from a database, and how to read and write serialized data to a file.

Pex Files database/windb.pex.

Mode Distributed only.

Special Requirements Database connection. The files artist.dat and painting.dat must be located in FORTE_ROOT/install/examples/database.

► **To use WinDB**

1. You need an environment that has a node with a resource manager.

Before running WinDB, open the WinDB Project Workshop, and open the MySession service object. Provide the correct values for your database in the Database Manager, Database Name, User Name, and User Password fields, then click OK. WinDB creates its own tables in the database you specified. The data is provided from files in the FORTE_ROOT/install/examples/database directory, called artist.dat and painting.dat.

2. Start the application. In the radio list, Table is selected by default. Click on the Make Database button. The following painter names are valid:
 - Leonardo da Vinci
 - Henri Rousseau
 - Edgar Degas
 - Jaspar Johns
 - Pablo Picasso
3. Enter a valid painter name and click the Select button. (Note that additional painting data is available for Edgar Degas and Leonardo da Vinci only.)
4. Enter a valid first letter of a painter name followed by% (such as E%), and click the Select button.
5. Enter an invalid painter name and click the Select button.
6. Click on the Drop Database button.
7. Now select Blob from the radio list, click on the Make Database button, and make the same selections as you did for Table. Click the Drop Database button when you're done.
8. Now select File from the radio list, and click on the Make Database button. This will actually create a file to which objects are written. This time, when you select artists, you must type their entire name. Again, click the Drop Database button when you're done.

TOOL SQL Statement Reference

This appendix is the primary reference for the TOOL SQL statements:

- sql close cursor
- sql delete
- sql execute immediate
- sql execute procedure
- sql fetch cursor
- sql insert
- sql open cursor
- sql select
- sql update

These statements are a subset of all TOOL statements. The *TOOL Reference Guide* provides complete reference regarding TOOL language components and usage, and all other TOOL statements.

In general, any operation you can perform using TOOL SQL you can also perform using a GenericDBMS class. However, for all databases except Sybase, you must use the TOOL SQL statement execute procedure to execute a database procedure; you cannot execute procedures using classes in GenericDBMS.

Note on Vendor-Specific SQL Extensions

You can include vendor-specific extensions to the ANSI SQL syntax in your TOOL SQL statements. However, if you include vendor-specific clauses in TOOL statements, your code will be valid only for database systems that support that particular syntax extension.

You can make vendor-specific code conditional upon the type of database for the current database session; see [“Using Conditional TOOL for Vendor-Specific Code” on page 96](#) for information on using vendor-specific extensions in the iPlanet UDS SQL statements.

A small number of vendor-specific database features are explicitly not supported; see [“Unsupported Database Features” on page 35](#).

Because the iPlanet UDS parser does not explicitly recognize vendor extensions, it cannot verify whether they are syntactically or semantically correct. Also, TOOL code that compiles may run without exception against one database but not against another, due to different support for the SQL language.

SQL Close Cursor

The `sql close cursor` statement closes a cursor that was opened with the `sql open cursor` statement.

Syntax

```
sql close cursor cursor_reference;
```

Example

```
sql close cursor dbcursor;
```

Description

After a cursor is closed, you cannot use it again until you open it. If you close a cursor before all rows have been fetched, the remaining rows are not fetched.

To specify the cursor you wish to close, you must specify a currently open cursor using a cursor reference (see [“Cursor Reference” on page 184](#)).

SQL Delete

The `sql delete` statement removes rows from a table.

Syntax

```
sql delete from table_name
    [where {search_condition | current of cursor_reference}]
    [on session {session_object_reference | default}];
```

or

```
{numeric_attribute | numeric_variable} = (sql_delete_statement);
```

Example

```
sql delete from artist_table where born < 1500;
```

Description

The `sql delete` statement removes the specified rows from a database table. The `where` clause identifies the particular rows to be deleted. If you do not include the `where` clause, all rows in the table are deleted.

Return Value

If the `sql delete` statement succeeds, it returns a numeric value indicating the number of rows that were deleted. A value of 0 indicates that no rows were deleted because a matching row was not found. A value of 1 or more indicates the number of rows that were deleted.

To assign the return value to a numeric variable or attribute, you must enclose the entire `sql delete` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in numeric expressions.

If the `sql delete` statement fails, iPlanet UDS raises an exception. The return value does not indicate whether there was an error.

Table Name

The table name identifies the database table that contains the rows to be deleted. The table must be available to the database session used for the `sql delete` statement. See “On Session Clause” below.

Where Clause

The optional `where` clause specifies the rows to be deleted. If you omit the `where` clause, all rows in the table are deleted. If you include a `where` clause, you use either a search condition, or, if your particular database system permits it, a cursor.

The search condition for the `where` clause specifies a condition that deleted rows must meet. If the condition is TRUE for a row, the row is deleted. If the condition is FALSE for a row, the row is retained.

The search condition is the same as the search condition in the `where` clause of the `sql select` statement (see “SQL Select” on page 185). If you use an iPlanet UDS variable or attribute in the `where` clause, you must preface it with a colon to distinguish it from a database name. The following example illustrates:

```
name : TextData = new(value = 'Leonardo da Vinci');
sql delete from artist_table where name = :name;
```

For databases that allow positioned update, you can use a cursor with the `sql delete` statement to delete the single row to which the cursor is pointing (that is, the last row fetched). To identify the cursor, you must use a cursor reference. The cursor must be defined for the table specified in the `delete` statement and the cursor must include the `for-update` clause.

Before you can use a cursor with the `sql delete` statement, you must use the `sql open cursor` statement in an explicit transaction. Then you use the `sql fetch` statement to position the cursor on the row you want to delete. Finally, you use the `where current of` clause in the `sql delete` statement with a reference to that cursor. Note that when you use the `where current of` clause, you cannot use the `on session` clause, because iPlanet UDS uses the session associated with the cursor when the cursor was opened.

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

SQL Execute Immediate

The `sql execute immediate` statement executes one SQL statement.

Syntax

```
sql execute immediate
{string_literal | string_variable | attribute | TextData_reference}
[on session {session_reference | default}];
```

or

```
{numeric_attribute | numeric_variable} =
(sql_execute_immediate_statement);
```

Example

```
sql execute immediate 'create table xyz (col int)'
on session dbsess;
```

Description

The `sql execute immediate` statement executes a single SQL statement that you enter directly as a literal string or that is stored in an iPlanet UDS variable or attribute.

The `sql execute immediate` statement allows you to execute SQL statements that you cannot include directly in a method or that are constructed at runtime. For example, you can use `sql execute immediate` to execute a create table or create index statement.

The literal string can be any SQL statement that does not return results. Be sure to enclose the string in single quotes.

You can use either a variable or attribute of type string or TextData to specify the SQL statement. If you use a string, its value must equal a single SQL statement. If you use a TextData object, its Value attribute must be set to a single SQL statement.

The `sql execute immediate` statement can be used to issue data definition language (DDL) statements, such as `create table`. Database systems differ in how they treat DDL statements within a transaction. They may:

- permit DDL statements within multi-statement transactions
- prohibit DDL statements within multi-statement transactions
- cause an implicit commit of the transaction in progress.

The following table summarizes the behavior of each iPlanet UDS-supported DBMS:

Database	Interaction between a DDL statement and a transaction
DB2	Allows one or more DDL statements in a multi-statement transaction.
Informix	Allows one or more DDL statements in a multi-statement transaction.
Ingres	Allows one or more DDL statements in a multi-statement transaction.
ODBC	Uses the semantics of the underlying database.
Oracle	Use of a DDL statement automatically commits the current transaction.
Rdb	Allows one or more DDL statements in a multi-statement transaction.
Sybase	Allows only one DDL statement in a transaction.

If your application must be portable across databases, do not put `execute immediate` statements containing DDL in transactions that contain more than one SQL statement. Instead, enclose each DDL statement in its own explicit transaction block.

Return Value

If the `sql execute immediate` statement succeeds, it returns a numeric value indicating the number of rows that were affected. A value of 0 indicates that no rows were affected. A value of 1 or more indicates the number of rows that were affected.

To assign the return value to a numeric variable or attribute, you must enclose the entire `sql execute immediate` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in numeric expressions.

If the `sql execute immediate` statement fails, iPlanet UDS raises an exception. The return value does not indicate whether there was an error.

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

SQL Execute Procedure

The `sql execute procedure` statement executes a database procedure.

Syntax

```
sql execute procedure procedure_name
  [ ([ input | output | input output] parameter = expression
    [, [ input | output | input output] parameter = expression...] ) ]
  [on session {session_object_reference | default}];
```

or

```
DataValue_reference = ( sql_execute_procedure_statement );
```

Example

```
empid : integer = 12345;
salaryIncrement : integer = 15000;
-- Passing parameters by position.
sql execute procedure updateSalary(empid, salaryIncrement);
```

```
-- Passing parameters by name.  
sql execute procedure updateSalary(AddToSalary = salaryIncrement,  
    Id = empid);
```

Description

The `sql execute procedure` statement executes a database procedure using the parameter values you specify. If your database supports it, the procedure can use output parameters to return values to your iPlanet UDS application.

NOTE You must use this TOOL SQL statement to execute database procedures for all databases *except Sybase*; you cannot execute procedures using classes in the GenericDBMS Library. However, to process results from a Sybase stored procedure, you must use methods in the GenericDBMS classes; you cannot use this TOOL statement. See [“Executing Prepared Queries” on page 120](#) for more information.

Return Value

If the procedure produces a return value, you can assign the return value to an iPlanet UDS variable or attribute. The `sql execute procedure` statement returns a `DataValue` object reference, which is the return value for the procedure. The actual class of the `DataValue` object corresponds to the return type of the procedure.

The variable or attribute to which you assign the return value must be a `DataValue` class, and the type must be compatible with the data type of the return value (see [Chapter 4, “Working with Data Types”](#) for information about data type compatibility). When you assign the return value to an attribute or variable, you must enclose the `sql execute procedure` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in expressions.

Note that the return value for this statement does not indicate whether there was an error. If the statement fails, iPlanet UDS raises an exception.

Procedure Name

The procedure name identifies the database procedure to be executed. This can be any procedure available to the database session associated with the `execute procedure` statement (see [“On Session Clause”](#) below).

Parameter List

The optional parameter list provides input values for the procedure. If your procedure is defined with output or input-output parameters, you can also use this parameter list to specify variables or attributes to store the output values. Although the input, input-output, and output parameters for database procedures appear to be the same as the parameters for iPlanet UDS methods, the following differences apply:

- the database contains and controls their definitions
- for database parameters, you must specify whether they are input or output at the point of call because iPlanet UDS does not have the parameter definition
- you must make sure that the call matches the original parameter definition because the iPlanet UDS compiler cannot check this

NOTE Do not use colons to identify the iPlanet UDS names in the parameter list in the `sql execute procedure` statement.

You must specify values for all required parameters (that is, parameters that do not have default values). You do not need to specify values for optional parameters (parameters having default values). Any parameter that you omit uses its default value.

To specify an input parameter value with an optional name, enter the name of the parameter as specified in the procedure declaration followed by the value. The value can be any expression that is compatible with the data type of the parameter (see [Chapter 4, “Working with Data Types”](#) for information about data type compatibility). The `input` key word before the parameter name is not required, because this is the default.

To specify an output or input-output parameter value, use the `output` or `input output` options to identify the kind of parameter. Then enter the name of the parameter as specified in the procedure declaration followed by the iPlanet UDS variable or attribute to specify the input value (for input-output parameters) and contain the output value (for both input-output and output parameters). The data type of the variable can be any simple type or `DataValue` type that is compatible with the data type of the parameter.

Because iPlanet UDS cannot verify the parameter names or types at compile time, you get a runtime error if they are incorrect. You will also get a runtime error if you try to use the `input output` or `output` options and the parameter was not originally declared for that purpose.

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

Exceptions

Any subclass of `DBResourceException`.

SQL Fetch Cursor

The `sql fetch cursor` statement allows you to fetch the next row, the next *n* rows, or the entire result set.

Syntax

```
sql fetch [[next {integer_constant | :integer_varName}] from]
cursor cursor_reference
[into {object_reference | simple_list | array_reference}];
or
{numeric_variable | numeric_attribute} = (sql_fetch_cursor_statement);
```

Example

```
sql fetch cursor dbcursor into :a;
```

Description

If you fetch one row, `sql fetch cursor` moves the cursor to the next row in the result set and assigns the values in the row to the specified attributes or variables.

If you fetch multiple rows, the `sql fetch cursor` statement fetches either the entire result set or the specified number of rows into the iPlanet UDS array specified in the `into` clause. Fetching multiple rows is more efficient than fetching rows one at a time. In addition, if array support exists in the DBMS, iPlanet UDS takes advantage of this.

Before you can use `sql fetch`, you must first open the cursor using the `sql open cursor` statement. After you open the cursor, the cursor is positioned before the first row in the result set.

If you are fetching rows one at a time, you can use the `sql fetch` statement in a loop to process each row in the result set. The first time you use the `sql fetch` statement, iPlanet UDS moves the cursor to the first row in the result set. Then with each successive `sql fetch` statement, iPlanet UDS moves the cursor forward one row, until it fetches the last row in the result set. An example follows:

```

-- Assume that the following cursor is declared.
cursor artist_cursor (name : string)
begin
    select name, country, born, died, school, comments
    from artist_table
    where name LIKE :name;
end;
-----
-- The following code will use the cursor.
dbsess : DBSession = ... more here ...
dbcursor : artist_cursor;
a : Artist = new;
painters : Array of Artist = new;
sql open cursor dbcursor ('%') on session dbsess;
while (sql fetch cursor dbcursor into :a)
    > 0 do
    -- Since 'a' will be reused, make a new one after append.
    painters.AppendRow(a);
    a = new;
end while;
sql close cursor dbcursor on session dbsess;
-- Now print them out
for x in painters do
    x.WriteToLog();
end for;

```

If you use the optional key word `next`, you can specify that a given number of rows should be fetched, and you must use an array reference in the `into` clause.

If the value following the next key word (say n) evaluates to a positive number, at most n rows are fetched from the cursor. If less than n rows remain, they are all fetched. If n evaluates to 0, all remaining rows are fetched from the cursor. If n evaluates to less than 0, no rows are fetched. You can use the `Array.Items` attribute to find out how many rows were fetched and put into the array.

An example showing the use of the next keyword follows:

```

dbsess  : DBSession;
rows    : integer = 1;
dbcursor : artist_cursor;
painters : array of artist = new;

sql open cursor dbcursor ('%') ON SESSION dbsess;
while (rows > 0) do
    sql fetch next 10 from cursor dbcursor into :painters;
    rows = painters.Items;
    -- do something with painters ...
end while;

```

The `sql fetch cursor` statement is never enclosed in an implicit transaction. The behavior of a `sql fetch cursor` statement executed outside of an explicit `TOOL` transaction depends on when the cursor was opened. If the cursor was opened in an explicit `TOOL` transaction, then it was closed at the end of the transaction and any fetch outside the transaction in which it was opened will fail. If the cursor was opened outside of an explicit transaction, the result set for the cursor was buffered and the `fetch` statement will fetch the next row from the buffer.

Return Value

The `sql fetch cursor` statement returns a numeric value indicating the number of rows that were fetched. A return value of one indicates that iPlanet UDS successfully fetched one row. A return value of greater than one indicates that you fetched the entire result set into an array, and the number indicates the total number of rows in the array. A return value of zero indicates that there are no more rows left in the result set.

To assign this return value to a numeric variable or attribute, you must enclose the entire `sql fetch cursor` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in numeric expressions.

Note that the return value for this statement does not indicate whether there was an error. If the statement fails, iPlanet UDS raises an exception.

Cursor

To identify the cursor, you must use the cursor reference. This must be a cursor that you have opened with the `sql open cursor` statement.

Into Clause

The `into` clause identifies the attributes and/or variables in which you want to store the column values. You can specify either a single object reference, or a list of simple data items or `DataValue` objects. Note that when you use iPlanet UDS names in a SQL statement, you must precede each name with a colon.

You can reference an existing object in which to store the column values. The attributes in the object must correspond by name with the columns in the result set. They must also have compatible data types. If there are any attributes in the object that have class types (other than `DataValue` types), iPlanet UDS ignores these, even if they have matching names.

If there are more columns in the row than attributes in the object, the extra columns are ignored (you cannot access them from `TOOL`). If there are more attributes in the object than columns in the row, the extra attributes retain the values they had before the `sql fetch` statement.

The object must already have been created, and any attributes that are `DataValue` objects must be created before the `sql fetch` statement.

Note that if you wish to fetch the values from a table into an object whose attribute names do not match the column names, you can use the column list of the `select` statement for the cursor to specify correlation names for the columns.

If the object you reference is an array object, the `sql fetch` statement fetches either the entire result set in the cursor or the number of rows specified in the `next` clause into the array. If the array reference is `NIL`, iPlanet UDS creates a new array object and adds the rows in the result set to the array. If the array reference is for an existing array object, iPlanet UDS replaces the values of the attributes in each of the array rows with the values from the result set. If there are more rows in the result set than in the original array, iPlanet UDS adds the remaining rows as new rows. If there are more rows in the array, iPlanet UDS deletes the extra array rows.

The list of simple attributes and/or variables can include any attributes or variables with a simple data type or with a subclass of the `DataValue` class (such as `TextData`, `IntegerNullable`, or `DateTimeData`). The order of the attributes and/or variables must match the order of the columns in the result set, and they must have compatible data types (see [Chapter 4, “Working with Data Types”](#) for information on data type compatibility). The syntax is:

variable | attribute [, *variable | attribute*]..

If you include `DataValue` objects on the list of simple data items, you must create them before using the `sql fetch` statement.

In order to return null values from the database, you must assign the value to a nullable `DataValue` object. If you do not use a nullable `DataValue` object, you will get a runtime exception. Use the `IsNull` method to check if there is a null value. See the iPlanet UDS online Help for information on the nullable `DataValue` classes.

SQL Insert

The `sql insert` statement adds a new row to a database table.

Syntax

```
sql insert into table_name [(column [, column]...)]
  {values ({object_reference | simple_list)} | select_statement}
  [on session {session_object_reference | default }];
```

or

```
{numeric_attribute | numeric_variable} = (sql_insert_statement);
```

Example

```
a : Artist = new ... fill in data...;
sql insert into artist_table
  (name, country)
  values (:a.name, :a.country)
  on session dbsess;
```

Description

The `sql insert` statement adds a new row to the specified table. You can either provide a list of values to use for the new row or you can use a `select` statement to get the values from another table. If you wish to add a set of rows stored in an iPlanet UDS array, you can specify an array object in the `values` clause.

Return Value

The `sql insert` statement returns a numeric value indicating the number of rows that were inserted. A value of zero indicates that no rows were inserted. A value of one indicates that one row was inserted. If you are inserting an set of rows stored in an iPlanet UDS array, the return value indicates the total number of rows added to the table.

To assign this return value to any numeric variable or attribute, you must enclose the entire `sql insert` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in numeric expressions.

Note that the return value for this statement does not indicate whether there was an error. If the statement fails, iPlanet UDS raises an exception.

Table Name

The table name identifies the database table to which you wish to add a new row. This can be any table available from the database session that you are using with the `sql insert` statement.

The Column List

The column list identifies the columns for which you are providing values. Any columns that you do not include on this list will be set to their default values. If you are specifying the insert values with a list of simple variables and attributes, or with a `select` statement, the column list must match *by position* with the insert values. If you are specifying the insert values with an object reference, the column list must match the attributes in the object *by name*.

Specifying the Insert Values

You can specify the values for the new row either by using the `values` clause to enter a list of values or by using an `sql select` statement to select the values from another table.

The `values` clause provides the values to use for the new row. You can use a single object reference or a list of variables and attributes.

You can reference an existing object to provide the column values. The attributes in the object must correspond by name with the columns in the table (or the columns you specified in the column list). They must also have compatible data types. Any attributes that do not have simple data types or `DataValue` class types are ignored.

If you provided a column list for the statement, any attribute names that are not also on the column list will be ignored.

By specifying an array object in the `values` clause, you can insert an entire set of rows with a single `sql insert` statement. iPlanet UDS also uses any underlying optimizations that may be provided by a specific vendor. To match the attributes in the array rows with the columns in the table, iPlanet UDS uses name matching as described above for a single object.

The list of simple attributes and/or variables can include any attributes or variables with a simple data type or with a subclass of the `DataValue` class (such as the `TextData`, `IntegerNullable`, or `DateTimeData` classes). The order of the attributes and/or variables must be in the same order as the columns in the column list of the `sql insert` statement, and they must have compatible data types (see [Chapter 4, “Working with Data Types”](#) for information on data type compatibility). The syntax is:

```
variable | attribute [ , variable | attribute ] . . .
```

The `select` statement selects a row from another table and uses its values for the new row. You can use any valid `select` statement, without an `into` clause. The columns in the `select` statement’s column list must be in the same order as the columns in the table you are inserting into. If you specify an asterisk instead of a column list, the columns in the table you are selecting from must be in the same order as the table you are inserting into.

If the `select` statement selects more than one row, the `sql insert` statement uses the values from the first row.

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

SQL Open Cursor

The `sql open cursor` statement selects rows from a database table to be used with a database cursor.

Syntax

```
sql open cursor cursor_reference
    [ (expression [, expression]...) ]
    [on session {session_object_reference | default}];
```

Example

```
dbcursor : empcursor;
empid : integer;
sql open cursor dbcursor (empid) on session dbsess;
```

Description

The `sql open cursor` statement executes the `select` statement associated with the cursor and positions the cursor before the first row in the result set.

For an overview of defining and using cursors, see [“Selecting Multiple Rows using Cursors” on page 105](#).

When you define a cursor in the Cursor Workshop, you may define it as a *read-only* or *for-update* cursor. The effect of using either type of cursor varies by database vendor. One use of a for-update cursor is to use it in a “positioned update.” A positioned update uses the `where current of` clause in the `sql delete` or `sql update` statements to update (or delete) the current row. Of the database vendors, Oracle, Ingres, Informix, and Rdb support positioned update of cursors; Sybase, DB2 and ODBC do not support positioned updates.

As an alternative to using cursors, you can use the `TOOL for` statement to work with a result set of a `select` statement. The `for` statement is a simpler way to work with ad hoc queries; a cursor is more efficient for queries that will be re-executed, and the cursor can be shared. See the *TOOL Reference Guide* for information on the `for` statement.

For best performance, you should start an explicit iPlanet UDS transaction, using the `begin transaction` `TOOL` statement, before opening a cursor. While you can use `sql open cursor` outside of an explicit transaction, it is not recommended, as explained below.

An `sql open cursor` statement executed outside an explicit iPlanet UDS transaction is enclosed in an implicit transaction. Since the implicit transaction ends before the data is actually fetched, and because most DBMSs do not allow cursors to remain open when a transaction ends, the session automatically buffers the data as part of executing the `sql open cursor` statement. This means that the entire result set is retrieved for the cursor. Subsequent `sql fetch cursor` statements retrieve results from the cursor as usual, but the data comes from the buffer rather than from the database.

An implicit transaction for an open cursor buffers all the data and then automatically commits the transaction in the database. The actual data in the database corresponding to the buffered data can be concurrently updated without any notification to your program nor any locking on the data. Concurrent updating does not pose a problem for read-only data; however, to assure maximum consistency for data that will be updated, you should use explicit transaction statements.

Cursor Reference

When you work with cursors, you use a cursor reference. After a cursor is defined in the Cursor Workshop it has a cursor name. However, to reference a cursor in a TOOL SQL statement, you declare a cursor reference variable using the following format:

variable_name : *cursor_name*

Subsequently, in SQL statements that use that same cursor, you use the cursor reference.

Placeholder Assignment

The placeholder assignment list specifies the values for the placeholders used in the original cursor declaration. To specify these values, simply enter a list of values. The order of the values must match the order specified in the original cursor definition by position. The value for a placeholder can be any expression that is compatible with the data type of the placeholder.

NOTE Do not use colons to identify the iPlanet UDS names in the placeholder list in the `sql open cursor` statement.

In the following example, colons do not precede the iPlanet UDS variables `empid`, `ival`, and `fval`.

```
sql open cursor dbcursor (empid) on session dbsess;
integer ival = 4;
float fval = 5.2;
sql open cursor mycursor (ival, fval) on session dbsess;
```

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

SQL Select

The `sql select` statement retrieves one or more rows from one or more database tables.

Syntax

```
sql select [all | distinct] { * | column_list }
[into {object_reference | simple_list}]
[from table_name [, table_name]...]
[where search_condition]
[group by column_name [, column_name]...]
[having search_expression]
[order by column [asc | desc] [, column [asc | desc] ]... ]
[on session {session_object_reference | default}];
```

or

```
{numeric_attribute | numeric_variable} = (sql_select_statement);
```

Example

```
sql select * into :artist_object from artist_table
where name = :vname
on session dbss;
i : integer;
vname, vcountry : TextData = new;
i = (sql select name, country into :vname, :vcountry from
    artist_table on session dbss);
```

Description

The `sql select` statement allows you to retrieve rows from one or more database tables, and to store the values in iPlanet UDS variables or attributes.

You can use this statement to select a single row into an iPlanet UDS object or set of simple variables, or to select a set of rows to insert into an iPlanet UDS array. You can also use the `sql select` statement in a `for` statement and in a cursor declaration. (See the *TOOL Reference Guide* for information on using the `sql select` statement in a `for` statement. See [“Repeating a Statement Block” on page 108](#) for information about using `sql select` in the `cursor` statement.)

Unless you are selecting into an array object (see [“Into Clause” on page 188](#)), if you use `sql select` on its own, you can select only one row. If the result is more than one row, no data is returned and iPlanet UDS raises a `DBOperationException`.

Note that this section does not contain an exhaustive list of the valid syntax for every clause of the `select` statement for all databases. Generally speaking, iPlanet UDS supports almost all standard and vendor-specific supporting clauses, such as `group by`, `compute`, and so on.

Return Value

The `sql select` statement returns a numeric value indicating the number of rows that were selected. A value of zero indicates that there was no matching row. A value of one or greater indicates the total number of matching rows.

To assign this return value to any numeric variable or attribute, you must enclose the entire `sql select` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in numeric expressions.

Note that the return value for this statement does not indicate whether there was an error. If the statement fails, iPlanet UDS raises an exception.

Eliminating Duplicate Rows

When you are using a `for` loop or a cursor to select a set of rows, or if you are selecting into an array, you can use the `distinct` option to eliminate duplicate rows from the result set. This option ensures that the statement does not return identical rows. The `all` option specifies the default, which is that the statement returns all rows that match the `where` clause.

Column List

You must specify which columns you wish to retrieve. You can use the asterisk (*) to specify all columns or you can provide a list of column names. We recommend that you do specify the column names, even if you want to retrieve them all. iPlanet UDS can only optimize the `select` statement when the column names are specified at compile time.

The syntax for the column list is:

```
column [alias] [, column [alias]]...
```

If you specify a list of simple attributes and/or variables in the `into` clause (see below), these must match the columns in the column list by position. If you specify an object in the `into` clause, the attributes in the object must match the columns in the column list by name. However, if the object you wish to store the values in has attributes with different names than the columns, you can use the optional alias in the column list to “rename” the column.

In the following example, the `Employee` class has the following attributes: `empId`, `EmpFname`, and `empLname`. The columns in the database table have the following names: `id`, `firstName`, and `lastName`. By using aliases in the column list, the `sql select` statement provides the matching names to associate each column with the appropriate attribute in the object.

```
sql select id empId, firstName empFname,
        lastName empLname from empTable
into :empObject
where id = :id;
```

Note that the use of column renaming is subject to the restrictions of your database. For example, in `Rdb` you rename columns using the ANSI SQL “`AS`” clause.

Using aggregates and functions You can include aggregates and functions in the column list. iPlanet UDS lets you use any SQL construct that is supported by your particular DBMS. However, because aggregates and functions vary between databases, this may mean that your SQL is not portable.

Into Clause

The `into` clause identifies the attributes and/or variables in which you want to store the column values. You can specify a single object or a list of simple data items.

You can reference an existing object in which to store the column values. The attributes in the object must correspond by name with the columns in the column list (or the columns in the table if you use an asterisk). They must also have compatible data types (see [Chapter 4, “Working with Data Types”](#) for information on data type compatibility). If there are any attributes in the object whose type is a class (other than `DataValue` types), iPlanet UDS ignores these, even if they have matching names.

If there are more columns in the column list than attributes in the object, the extra columns are ignored (you cannot access them from `TOOL`). If there are more attributes in the object than columns in the column list, the extra attributes retain the values they had before the `sql select` statement.

The object must already have been created, and any attributes that are `DataValue` objects must be created before the `sql select` statement.

If the object you reference is an array object, the `sql select` statement fetches the entire result set into the array. If the array reference is `NIL`, iPlanet UDS creates a new array object and adds the rows in the result set to the array. If the array reference is for an existing array object, iPlanet UDS replaces the values of the attributes in each of the array rows with the values from the result set. If there are more rows in the result set than in the original array, iPlanet UDS adds the remaining rows as new rows. If there are more rows in the array than in the result set, iPlanet UDS deletes the extra rows.

The list of simple attributes and/or variables can include any attributes or variables with a simple data type or with a subclass of the `DataValue` class (such as `TextData`, `IntegerNullable`, or `DateTimeData`). The order of the attributes and/or variables must match the order of the columns in the result set, and they must have compatible data types (see [Chapter 4, “Working with Data Types”](#) for information on data type compatibility). The syntax is:

```
variable | attribute [, variable | attribute] . . .
```

If you include `DataValue` objects on the list of simple data items, you must create them before using the `select` statement.

From Clause

The `from` clause identifies one or more database tables from which you wish to select rows. You can include any tables available to the database session that you are using with the `sql select` statement.

Where Clause

The `where` clause specifies one or more conditions that the selected row or rows must meet. This can be any search condition that is allowed by your database management system.

You can use iPlanet UDS attributes and variables in the search condition to specify values. These are allowed anywhere that a literal is allowed. Be sure to precede all iPlanet UDS names with colons to distinguish them from column names.

You can include any SQL functions in your search condition that are allowed by your particular database. However, since supported functions vary among database systems, this may mean that your SQL is not portable.

If you omit the `where` clause, the `select` statement returns all the rows in the table.

Group By Clause

When you are using aggregates in the column list, you can use the `group by` clause to break the result set into groups. iPlanet UDS produces a summary value for each group rather than for the result set as a whole.

The `group by` clause specifies the columns to use as the basis for creating the groups. You can enter any columns from the `select` statement's column list. iPlanet UDS starts a new group each time the value changes in the specified column.

Having Clause

The `having` clause is for use with the `group by` clause. It specifies a condition that rows must meet in order to be included in the groups controlled by the `group by` clause. The search condition for the `having` clause is exactly the same as the search condition for the `where` clause, except that it can also include aggregates.

Order By Clause

The `order by` clause sorts the result set by one or more columns. You can specify any columns from the `select` statement's column list.

For each sort column you can specify an ascending sort or descending sort. If you do not specify either, the default is ascending.

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

Exceptions

Any subclass of `DBResourceException`.

SQL Update

The `sql update` statement changes values in one or more rows from a database table.

Syntax

```
sql update table_name set column = expr [, column = expr]...
    [where {search_expression | current of cursor_reference} ]
    [on session {session_object_reference | default}];
```

or

```
{numeric_attribute | numeric_variable} = (sql_update_statement);
```

Example

```
sql update artist_table set born = :vborn
    where name = :vname on session dbsess;
```

Description

The `sql update` statement replaces the current column values in the selected row or rows with the new values that you specify. You use the `set` clause to specify the new values for the columns. You use the `where` clause to select the rows that you wish to update. Without the `where` clause, the `sql update` statement changes the specified column values in all the rows in the table.

Return Value

The `sql update` statement returns a numeric value indicating the number of rows that were updated. A value of zero indicates that no rows were updated. A value of one or more specifies the total number of rows that were updated.

To assign this return value to any numeric variable or attribute, you must enclose the entire `sql update` statement in parentheses. See the *TOOL Reference Guide* for information about using SQL statements in numeric expressions.

Note that the return value for this statement does not indicate whether there was an error. If the statement fails, iPlanet UDS raises an exception.

Table Name

The table name identifies the database table that you wish to update. This can be any table from the database session that you are using with the `sql update` statement (see [“On Session Clause” on page 192](#)).

Set Clause

To specify the new values for the rows, you enter a list of column names with their corresponding values.

The value for a column can be any expression that is compatible with the data type of the column (see [Chapter 4, “Working with Data Types”](#) for information on data type compatibility). If you use iPlanet UDS names in the expressions, be sure to precede the names with colons.

Where Clause

The `where` clause selects the rows to be updated. You can specify a search condition or, if your particular database system permits it, a cursor. If you omit the `where` clause, all the rows in the table are updated.

The search condition for the `where` clause specifies a condition that the rows must meet. iPlanet UDS tests all the rows in the table. If the condition is TRUE for a row, the row is updated. If the condition is FALSE for a row, the row is not updated.

The search condition is the same as the search condition in the `where` clause of the `sql select` statement (see “Where Clause” on page 189). If you use iPlanet UDS variable or attribute names in the `where` clause, be sure to preface them with colons. This distinguishes the iPlanet UDS names from the column names.

```
name : string = 'Leonardo da Vinci';
whenborn : integer = 1922;
sql update artist_table set born = :whenborn
  where name = :name on session dbsess;
```

For databases that allow positioned update, you can use a cursor with the `sql update` statement to update the single row to which the cursor is pointing (that is, the last row fetched). To identify the cursor, you must use a cursor reference. The cursor must be defined for the table specified in the `update` statement, it must include the `for-update` clause, and if updating is restricted to certain columns, you can update only those columns.

Before you can use a cursor with the `sql delete` statement, you must use the `sql open cursor` statement in an explicit transaction. Then you use the `sql fetch` statement to position the cursor on the row you want to update. Finally, you use the `where current of` clause in the `sql delete` statement with a reference to that cursor. Note that when you use the `where current of` clause, you cannot use the `on session` clause, because iPlanet UDS uses the session associated with the cursor when the cursor was opened.

On Session Clause

The `on session` clause identifies a database session to use for the execution of this statement. You can use a service object name or a reference to an object of the `DBSession` class.

The `default` keyword specifies the `DBSession` object associated with a variable or attribute named `DefaultDBSession`. If you omit the `on session` clause, it is the same as specifying `default`.

A

- AddExternalRM Escript command 44
- Array interface support 33
- Arrays
 - fetching into 108
 - selecting data into 99, 104
 - SQL insert statement with 182
 - SQL select statement with 188
- Autocommit 154

B

- Binary data
 - Informix 113
 - inserting 114
 - selecting 113
 - using 113
- BinaryData objects
 - deserializing 113
- BLOB data, inserting 114

C

- Casting 89
- Class types, user-defined 102
- ClassName for database vendor 69

- CloseCursor method
 - described 118
 - using 127
- CloseExtent method
 - described 118
- Columns
 - case sensitive names 109
 - mapping attribute names to 100
 - mapping to database 80, 94
 - pseudo 102
- Conditional code 96
- Connect method
 - connection options 63
 - dynamic connection 68
- ConnectDB method
 - connection options 63
 - and DBSession object 56
 - example 67
- Connecting to the database 55
- Creating a service object 56
- Creating objects
 - for statement 99
 - SQL fetch cursor statement 99
 - SQL select 98
 - SQL select into array 99
- Cursor reference 184
- Cursor workshop 105
- Cursors
 - example definition 106
 - hold 35
 - opening 125
 - read-only, for-update 183

Cursors (*continued*)

- reference vs. name 106
- retrieving rows 107
- scroll 35
- SQL close cursor statement 168
- SQL fetch cursor statement 176
- SQL open cursor statement 182
- using 105

D

Data

- accessing 79, 93
- inserting 110
- nullable 80
- overflow 81, 82

Data source, ODBC

- definition 36
- format 51
- name 61

Data type conversion

- DB2 84
- NULL values 81
- ODBC data sources 87
- Oracle 88
- overview 80–83
- Rdb 89
- Sybase 92

Data types

- DataValue class 80
- simple TOOL 80

Database connections

- connecting 55
- disconnecting 76
- options 62

Database Integrator (DBI) 34

Database name

- compared to resource name 43
- logical 47
- specifying at runtime 68
- vendor-specific format 60

Database procedures. *See* Procedures, database

Database sessions

- DBSession object 56
- default 73
- ending 76
- multiple 150
- multitasking 150
- starting 55, 65
- with no service object 68

Database vendors, specifying at runtime 69

DataValue class, nullable variants 81

DB2

- data type conversion 84
- db2cschrc file 45
- DB2INSTANCE environment variable 45, 61
- db2profile file 45
- library name 69
- pseudo-columns 103
- resource manager 32
- ResourceName parameter format 61

db2cschrc file 45

DB2INSTANCE environment variable 45, 46

db2profile file 45

DB2-specific information 46, 154

DBDataSet class

- placeholders 130

DBDeadlockException class 150

DBResourceMgr service objects

- definition 57
- using 66

DBSession class

- using for runtime SQL 118

DBSession object, database vendor specific 70

DBSession service objects

- creating 58
- default 73
- definition 57
- overriding login information 72
- using 65

DBStatementHandle class

- using 124

DBVendorType attribute

- using 97

DDL statements 142

Deadlocks 149
 DefaultDBSession service object 73
 Deleting a row 111
 Disconnect method 76
 Disconnecting from the database 76
 DynamicDataAccess sample application 163
 DynamicSQL sample application 164

E

Encina 153
 Environment variables
 for resource manager 44
 for user name and password 71
 iPlanet UDS-specific 45
 LIBPATH 44
 Error handling
 DBRemoteAccessException class 158
 Errors, database 157
 Execute method
 described 118
 using 131
 using placeholders 130
 ExecuteImmediate method
 described 119
 using 119
 Executing
 database procedures 111
 ExecuteImmediate method 119
 insert, update, delete statements 127
 single SQL statements 111
 Explicit transactions
 Informix scroll cursor 136
 recommendation 132
 sql fetch statement 107
 using 140
 ExtendCursor method
 described 119

F

FetchCursor method
 described 119
 scroll cursor 136
 using 126
 FileKey for database vendor 69
 FindLibrary method for database vendor 69
 For update clause 106, 170, 192
 FORTE_DB_MAX_STATEMENTS environment variable 45
 FORTE_DB_VENDORFLG environment variable 45
 FORTE_STARTUP 47
 fortedef.csh file 44
 fortedef.sh file 44

G

GenericDBMS classes
 compared to TOOL SQL 94

H

Hold cursors 35

I

ImageData class
 formats 113
 inserting into database 113
 selecting from database 113
 Informix
 hold cursors 35
 INFORMIXDIR environment variable 45, 46
 INFORMIXSERVER environment variable 45, 46, 61

Informix (*continued*)

- library name 69
- pseudo-columns 103
- resource manager 33
- resourceName parameter format 61
- scroll cursors 35, 136

Informix-specific information 46, 76, 112, 154

Inserting data

- binary 114
- from objects 109
- from variables 109
- multiple rows 110
- single row 109
- using a column list 110

Instantiating objects 110

Interfaces file (Sybase) 47

IsConnected attribute 75

L

LIBPATH environment variable 44

M

Mutex lock 148

N

Names

- columns and attributes 101
- iPlanet UDS 95
- placeholders 130

Node properties dialog 41

NULL data, transferring 80

O

Objects. *See* Creating objects; Selecting data;
Transactional objects

ODBC

- architecture 35
- array interface support requirements 37
- data source 35, 51
- data type conversion 86
- drivers 35
- flat file data source 54
- library name 69
- resource manager 53
- resourceName parameter format 61
- SQL Extended Fetch feature 37
- SQL Param Option feature 37

ODBC-specific information 53, 112

OpenCursor method

- described 119
- using 125
- using placeholders 130

Oracle

- column aliases 103
- data type conversion 88
- library name 69
- OpenVMS usage notes 77
- ORACLE_SID 61
- resource manager 33
- ResourceName parameter format 61

ORACLE_HOME 46

ORACLE_SID 47

Oracle-specific information 77

P

Partitioning 64

PDF files, viewing and searching 22

Placeholders

- assigning values 184
- DBDataSet class 130
- definition 130
- in prepared statements 120, 127

- Positioned update 106
 - SQL delete 170
 - SQL update 192
 - vendor support for 183
- Prepare method
 - described 119
 - scroll cursor 136
 - using 124, 129
- Prepared statements
 - executing 120
 - select 120
- PreparePositioned method
 - described 119
- Preparing a select statement 124
- Procedures, database
 - executing 111

R

- Rdb
 - data type conversion 89
 - environment variables 47
 - library name 69
 - pseudo-columns 103
 - renaming columns 103
 - resource manager 34
 - resourceName parameter format 61
 - transactions 155
- Rdb-specific information 78, 155
- RDMS\$DEBUG_FLAGS 45, 47
- RDMS\$RUJ 45, 47
- ReadFromFile method 113
- Reconnect method
 - using 75
- RemoveStatement method
 - described 119
- Resource manager
 - and resource name 39
 - DB2 32
 - defining a 39
 - environment variables 44
 - Informix 33

- name 43
- name space 43
- ODBC 53
- Oracle 33
- Rdb 34
- removing 47
- Resource name
 - definition 40
- resourceName parameter format 61
- Result set, fetching rows from 126
- ResultSet parameter, using 125
- Retrieving rows
 - fetching into an array 108
 - SQL fetch cursor statement 107
 - using a for statement 108
- Rows
 - deleting 111
 - inserting multiple 110
 - inserting single 109
 - retrieving 107–109
 - updating 110

S

- Sample applications
 - DynamicDataAccess 163
 - DynamicSQL 164
 - WinDB 165
- Scroll cursors 35
- Select method
 - described 119
- Select statements
 - building dynamically 122
 - preparing 124
- Selecting data
 - a single row 100
 - attributes and column names differ 101
 - binary 113
 - with for statements 105
 - inherited attributes 104
 - into a variable 100
 - into an array 99
 - into an object 100

- Selecting data (*continued*)
 - matching attributes and columns 101
 - multiple rows into arrays 104
 - using cursors 105
- SORTWORK rdb logical name 45, 47
- SQL
 - DDL statements 142
 - using DBStatementHandle 124
 - using DBVendorType with 97
 - vendor-specific extensions 96, 168
- SQL as clause (column renaming) 103
- SQL close cursor statement 168
 - definition 27, 95
 - syntax 168
 - using 105
- SQL delete statement 169–170
 - definition 27, 95
 - on session clause 170
 - positioned update 183
 - return value 169
 - syntax 169
 - table name 170
 - using 111
 - where clause 170
- SQL execute immediate statement 171–173
 - definition 28, 96
 - exceptions 176
 - on session clause 173, 176
 - parameter list 175
 - procedure name 174
 - return value 172
 - syntax 171
 - using 111
- SQL execute procedure statement 173–176
 - definition 28, 96
 - return value 174
 - syntax 173
 - using 111
- SQL fetch cursor statement 176–180
 - cursor 179
 - definition 27, 95
 - into clause 179
 - return value 178
 - syntax 176
 - using 99, 105–108
- SQL insert statement 180–182
 - array of rows 182
 - column list 181
 - definition 27, 95
 - insert values 181
 - on session clause 182
 - return value 181
 - syntax 180
 - table name 181
 - using 109
- SQL open cursor statement 182–185
 - cursor_reference 184
 - definition 27, 95
 - on session clause 185
 - placeholder assignment 184
 - syntax 183
 - using 105–107
- SQL select statement 185–190
 - all option 187
 - definition 27, 95
 - distinct option 187
 - exceptions 190
 - from clause 189
 - group by clause 189
 - having clause 189
 - into clause 188
 - object creation 98
 - on session clause 190
 - order by clause 189
 - return value 186
 - selecting into an array 188
 - syntax 185
 - using 98, 100
 - where clause 189
- SQL update statement 190–192
 - definition 27, 95
 - on session clause 192
 - positioned update 183
 - return value 191
 - set clause 191
 - syntax 190
 - table name 191
 - using 110
 - where clause 191
- SQL*Net (Oracle), using 77

Sybase

- data type conversion 92
 - interfaces file 45, 47
 - library name 69
 - resourceName parameter format 61
 - stored procedures 174
- SYBASE environment variable 45
- Sybase-specific information 47

T

- tnsnames.ora file (Oracle SQL*Net) 46, 77
- TOOL SQL statements 93–98
- list of 95
 - vendor-specific extensions 96
- Transactional objects 147
- Transactions
- explicit 140
 - implicit 141
 - Informix 154
 - iPlanet UDS 140
 - SQL open cursor 143

U

- Updating a row 110
- User name
- DBSession service objects 60
 - dynamic 71
 - overriding 66
 - specifying at runtime 68
- User password 60

V

- Variable user names and passwords 71
- Variable, selecting data into 100

W

- Where current of clause 106, 183
- WinDB sample application 165

