



J/Foundation Developer's Guide



J/Foundation Developer's Guide

Note!

Before using this information and the product it supports, read the information in “Notices” on page B-1.

First Edition (December 2004)

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1996, 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	vii
About This Manual	vii
Types of Users	viii
Software Dependencies	viii
Assumptions About Your Locale	viii
Demonstration Database	ix
New Features in J/Foundation, Version 10.0	ix
New Features in J/Foundation, Version 9.3 and 9.4	ix
Documentation Conventions	x
Typographical Conventions	x
Feature, Product, and Platform	xi
Syntax Diagrams	xii
Example Code Conventions	xv
Additional Documentation	xvi
Installation Guides	xvi
Online Notes	xvi
Informix Error Messages	xviii
Manuals	xix
Online Help	xix
Accessibility	xix
IBM Informix Dynamic Server Version 10.0 and CSDK Version 2.90 Documentation Set	xix
Compliance with Industry Standards	xxii
IBM Welcomes Your Comments	xxiii
Chapter 1. Concepts	1-1
Features of Java User-Defined Routines	1-1
Java Virtual Processors	1-1
Thread Scheduling	1-2
Query Parallelization	1-2
System Catalog Tables	1-2
Chapter 2. Preparing for Java Support	2-1
Installing the JDBC Driver	2-1
Configuring Java Support	2-2
Creating an sbspace	2-2
Creating the JVP Properties File	2-2
Setting Configuration Parameters	2-3
Setting Environment Variables	2-4
GLS Support	2-5
Chapter 3. Configuration Parameters	3-1
JVPARGS	3-1
JDKVERSION	3-2
JVPCLASSPATH	3-2
JVPHOME	3-3

JVPJAVAHOME	3-3
JVPJAVALIB	3-4
JVPJAVAVM	3-4
JVPLOGFILE	3-5
JVPPROFILE	3-6
SBSPECNAME	3-6
VPCLASS JVP	3-6
Chapter 4. Creating Java User-Defined Routines	4-1
Java User-Defined Routines	4-2
Creating a Java User-Defined Routine	4-3
Writing a Java User-Defined Routine	4-4
The com.informix.udr Package	4-5
The com.informix.udr.UDRManager	4-5
The com.informix.udr.UDREnv	4-5
The com.informix.udr.UDRLog	4-7
The com.informix.udr.UDRTraceable	4-7
Creating UDT-to-Java Mappings	4-8
Registering Java User-Defined Routines	4-9
Specifying the JVP	4-9
Using Routine Modifiers	4-10
Specifying the External Name	4-11
Using a Deployment Descriptor	4-13
Using a Manifest File	4-14
Compiling the Java Code	4-14
Installing a JAR File	4-15
Updating JAR Filenames	4-16
Executing the User-Defined Routine	4-16
Debugging and Tracing	4-16
Generating Log Messages.	4-16
Using the Administrative Tool	4-17
Debugging a Java User-Defined Routine.	4-18
Traceable Events.	4-18
Finding Information about User-Defined Routines	4-18
Complying with SQLJ	4-19
Unsupported Modifiers	4-20
Unsupported Optional Modifiers	4-21
Chapter 5. The IBM Informix JDBC Driver	5-1
Public JDBC Interfaces	5-1
The com.informix.jdbc.IfxConnection	5-2
The com.informix.jdbc.IfxProtocol	5-2
The informix-direct Subprotocol	5-2
JDBC 1.0 API	5-3
JDBC 2.0.	5-4
Support for Opaque Data Types	5-5
Interfaces Updated for Java 2.0	5-5
An Example That Shows Query Results	5-5
Chapter 6. Using Opaque User-Defined Types	6-1

Using the SQLData Interface	6-1
Default Input/Output Routines	6-2
SQL Definitions for Default I/O User-Defined Routines	6-3
Informix Extensions to SQLInput and SQLOutput	6-3
The Circle Class Example	6-5
Overriding the Default I/O Methods	6-6
I/O Function Sets and Related Types	6-6
An Example That Overrides the Default I/O Methods	6-9
Limitations to Streams	6-15
Appendix. Accessibility	A-1
Notices	B-1
Index	X-1

Introduction

About This Manual	vii
Types of Users.	viii
Software Dependencies	viii
Assumptions About Your Locale.	viii
Demonstration Database	ix
New Features in J/Foundation, Version 10.0	ix
New Features in J/Foundation, Version 9.3 and 9.4	ix
Documentation Conventions.	x
Typographical Conventions	x
Feature, Product, and Platform.	xi
Syntax Diagrams	xii
How to Read a Command-Line Syntax Diagram	xiii
Keywords and Punctuation	xiv
Identifiers and Names	xv
Example Code Conventions	xv
Additional Documentation.	xvi
Installation Guides	xvi
Online Notes	xvi
Locating Online Notes	xvii
Online Notes Filenames	xviii
Informix Error Messages	xviii
Manuals	xix
Online Manuals	xix
Printed Manuals	xix
Online Help	xix
Accessibility	xix
IBM Informix Dynamic Server Version 10.0 and CSDK Version 2.90 Documentation Set	xix
Compliance with Industry Standards	xxii
IBM Welcomes Your Comments	xxiii

In This Introduction

This introduction provides an overview of the information in this manual and describes the conventions it uses.

About This Manual

This manual describes how to write user-defined routines (UDRs) in the Java programming language for IBM Informix Dynamic Server with J/Foundation. It also describes the infrastructure that enables you to run Java applications in the database server. It describes the Java classes, methods, and interfaces that allow you to access databases from within IBM Informix Dynamic Server with J/Foundation, rather than from a client application.

This section discusses the organization of the manual, the intended audience, and the associated software products that you must have to develop and use Java UDRs.

Types of Users

This manual is written for the following users:

- Database-application programmers
- DataBlade module developers
- Java UDR developers
- Java server application developers

This manual assumes that you have basic knowledge in the following areas:

- Your computer, your operating system, and the utilities that your operating system provides
- Object-relational databases or exposure to database concepts
- The Java language and the Java Developer's Kit
- Java Database Connectivity (JDBC), which is a Java application programming interface to SQL databases
- SQLJ: SQL Routines specification, which specifies the Java binding of SQL UDRs

If you have limited experience with object-relational databases, SQL, or your operating system, refer to *IBM Informix: Dynamic Server Getting Started Guide* for a list of supplementary titles.

Software Dependencies

This manual assumes that you are using the following software:

- IBM Informix Dynamic Server with J/Foundation
- The Java Development Kit (JDK), Version 1.4.x

You need the JDK to compile your Java programs. However, J/Foundation includes Version 1.4.x of the Java Runtime Environment (JRE) and uses it to execute your server-based Java routines. This specific version of the JRE ensures that the Java environment is known and reliable for this database server release.

In addition, the DataBlade Developer's Kit (DBDK) for Java facilitates DataBlade module development.

Assumptions About Your Locale

IBM Informix products can support many languages, cultures, and code sets. All the information related to character set, collation, and representation of numeric data, currency, date, and time is brought together in a single environment, called a Global Language Support (GLS) locale.

The examples in this manual are written with the assumption that you are using the default locale, **en_us.8859-1**. This locale supports U.S. English format conventions for date, time, and currency. In addition, this locale supports the ISO 8859-1 code set, which includes the ASCII code set plus many 8-bit characters such as é, è, and ñ.

If you plan to use nondefault characters in your data or your SQL identifiers, or if you want to conform to the nondefault collation rules of character data, you need to specify the appropriate nondefault locale.

For instructions on how to specify a nondefault locale, additional syntax, and other considerations related to GLS locales, see the *IBM Informix: GLS User's Guide*.

Demonstration Database

The DB–Access utility, which is provided with the Informix database server products, includes one or more of the following demonstration databases:

- The **stores_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. Many examples in IBM Informix manuals are based on the **stores_demo** database.
- The **sales_demo** database illustrates a dimensional schema for data warehousing applications. For conceptual information about dimensional data modeling, see the *IBM Informix: Database Design and Implementation Guide*.

For information about how to create and populate the demonstration databases, see the *IBM Informix: DB–Access User's Guide*. For descriptions of the databases and their contents, see the *IBM Informix: Guide to SQL Reference*.

The scripts that you use to install the demonstration databases reside in the **\$INFORMIXDIR/bin** directory on UNIX platforms and in the **%INFORMIXDIR%\bin** directory in Windows environments.

New Features in J/Foundation, Version 10.0

In J/Foundation, Version 10.0, the JRE/JDK has been upgraded to 1.4x and the JDBC portion of Krakatoa has been upgraded to IBM Informix JDBC Driver, Version 2.21.JC4.

New Features in J/Foundation, Version 9.3 and 9.4

There are no new features in J/Foundation, Version 9.4.

The following table provides information about the new feature for IBM Informix Dynamic Server with J/Foundation, Version 9.3, which this manual covers. If you are accessing the manual online, click a blue hyperlink to go to the desired page. For a description of all new features, see the *IBM Informix: Getting Started Guide*.

Version 9.3 includes the following improvement for support of Java UDRs.

New Features	Reference
Java Virtual Machine (JVM) 1.3 support	For more information, see “JDKVERSION” on page 3-2.

Documentation Conventions

This section describes the conventions that this manual uses. These conventions make it easier to gather information from this and other volumes in the documentation set.

The following conventions are discussed:

- Typographical conventions
- Other conventions
- Syntax diagrams
- Command-line conventions
- Example code conventions

Typographical Conventions

This manual uses the following conventions to introduce new terms, illustrate screen displays, describe command syntax, and so forth.

Convention	Meaning
KEYWORD	All primary elements in a programming language statement (keywords) appear in uppercase letters in a serif font.
<i>italics</i> <i>italics</i> <i>italics</i>	Within text, new terms and emphasized words appear in italics. Within syntax and code examples, variable values that you are to specify appear in italics.
boldface <i>boldface</i>	Names of program entities (such as classes, events, and tables), environment variables, file and pathnames, and interface elements (such as icons, menu items, and buttons) appear in boldface.
monospace <i>monospace</i>	Information that the product displays and information that you enter appear in a monospace typeface.
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.

Convention	Meaning
>	This symbol indicates a menu item. For example, “Choose Tools > Options ” means choose the Options item from the Tools menu.

Tip: When you are instructed to “enter” characters or to “execute” a command, immediately press RETURN after the entry. When you are instructed to “type” the text or to “press” other keys, no RETURN is required.

Feature, Product, and Platform

Feature, product, and platform markup identifies paragraphs that contain feature-specific, product-specific, or platform-specific information. Some examples of this markup follow:

Dynamic Server

Identifies information that is specific to IBM Informix Dynamic Server

End of Dynamic Server

Extended Parallel Server

Identifies information that is specific to IBM Informix Extended Parallel Server

End of Extended Parallel Server

UNIX Only

Identifies information that is specific to UNIX platforms

End of UNIX Only

Windows Only

Identifies information that is specific to the Windows environment

End of Windows Only

This markup can apply to one or more paragraphs within a section. When an entire section applies to a particular product or platform, this is noted as part of the heading text, for example:

Table Sorting (Linux Only)

Syntax Diagrams

This guide uses syntax diagrams built with the following components to describe the syntax for statements and all commands other than system-level commands.

Note: Starting in 2004, syntax diagrams have been reformatted to conform to the IBM standard.

Syntax diagrams depicting SQL and command-line statements have changed in the following ways:

- The symbols at the beginning and end of statements are now double arrows instead of a vertical line at the end.
- The symbols at the beginning and end of syntax segment diagrams are now vertical lines instead of arrows.
- How many times a loop can be repeated is now explained in a diagram footnote instead of a number in a gate symbol.
- Syntax statements that are longer than one line now continue on the next line instead of looping down with a continuous line.
- Product or condition-specific paths are now explained in diagram footnotes instead of icons.

The following table describes syntax diagram components.

Component represented in PDF	Component represented in HTML	Meaning
	>>-----	Statement begins.
	----->	Statement continues on next line.
	>-----	Statement continues from previous line.
	-----<<	Statement ends.
	-----SELECT-----	Required item.
	---+-----+--- '-----LOCAL-----'	Optional item.
	---+-----+--- +---DISTINCT-----+ '-----UNIQUE-----'	Required item with choice. One and only one item must be present.

shown in uppercase letters. When you use a keyword in a command, you can write it in uppercase or lowercase letters, but you must spell the keyword exactly as it appears in the syntax diagram.

You must also use any punctuation in your statements and commands exactly as shown in the syntax diagrams.

Identifiers and Names

Variables serve as placeholders for identifiers and names in the syntax diagrams and examples. You can replace a variable with an arbitrary name, identifier, or literal, depending on the context. Variables are also used to represent complex syntax elements that are expanded in additional syntax diagrams. When a variable appears in a syntax diagram, an example, or text, it is shown in *lowercase italic*.

The following syntax diagram uses variables to illustrate the general form of a simple SELECT statement.

►—SELECT—*column_name*—FROM—*table_name*—►

When you write a SELECT statement of this form, you replace the variables *column_name* and *table_name* with the name of a specific column and table.

Example Code Conventions

Examples of SQL code occur throughout this manual. Except as noted, the code is not specific to any single IBM Informix application development tool.

If only SQL statements are listed in the example, they are not delimited by semicolons. For instance, you might see the code in the following example:

```
CONNECT TO stores_demo
...

DELETE FROM customer
    WHERE customer_num = 121
...

COMMIT WORK
DISCONNECT CURRENT
```

To use this SQL code for a specific product, you must apply the syntax rules for that product. For example, if you are using DB–Access, you must delimit multiple statements with semicolons. If you are using an SQL API, you must use EXEC SQL at the start of each statement and a semicolon (or other appropriate delimiter) at the end of the statement.

Tip: Ellipsis points in a code example indicate that more code would be added in a full application, but it is not necessary to show it to describe the concept being discussed.

For detailed directions on using SQL statements for a particular application development tool or SQL API, see the manual for your product.

Additional Documentation

For additional information, refer to the following types of documentation:

- Installation guides
- Online notes
- Informix error messages
- Manuals
- Online help

Installation Guides

Installation guides are located in the **/doc** directory of the product CD or in the **/doc** directory of the product's compressed file if you downloaded it from the IBM Web site. Alternatively, you can obtain installation guides from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>.

Online Notes

The following sections describe the online files that supplement the information in this manual. Please examine these files before you begin using your IBM Informix product. They contain vital information about application and performance issues.

Online File	Description	Format
TOC Notes	The TOC (Table of Contents) notes file provides a comprehensive directory of hyperlinks to the release notes, the fixed and known defects file, and all the documentation notes files for individual manual titles.	HTML
Documentation Notes	The documentation notes file for each manual contains important information and corrections that supplement the information in the manual or information that was modified since publication.	HTML, text
Release Notes	The release notes file describes feature differences from earlier versions of IBM Informix products and how these differences might affect current products. For some products, this file also contains information about any known problems and their workarounds.	HTML, text
Machine Notes	(Non-Windows platforms only) The machine notes file describes any platform-specific actions that you must take to configure and use IBM Informix products on your computer.	text
Fixed and Known Defects File	This text file lists issues that have been identified with the current version. It also lists customer-reported defects that have been fixed in both the current version and in previous versions.	text

Locating Online Notes

Online notes are available from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>. Additionally you can locate these files before or after installation as described below.

Before Installation

All online notes are located in the `/doc` directory of the product CD. The easiest way to access the documentation notes, the release notes, and the fixed and known defects file is through the hyperlinks from the TOC notes file.

The machine notes file and the fixed and known defects file are only provided in text format.

After Installation

On UNIX platforms in the default locale, the documentation notes, release notes, and machine notes files appear under the `$INFORMIXDIR/release/en_us/0333` directory.

Dynamic Server

On Windows the documentation and release notes files appear in the **Informix** folder. To display this folder, choose **Start > Programs > IBM Informix Dynamic Server *version* > Documentation Notes** or **Release Notes** from the taskbar.

Machine notes do not apply to Windows platforms.

End of Dynamic Server

Online Notes Filenames

Online notes have the following file formats:

Online File	File Format	Examples
TOC Notes	<i>prod_os_tocnotes_version.html</i>	ids_win_tocnotes_10.0.html
Documentation Notes	<i>prod_bookname_docnotes_version.html/txt</i>	ids_hpl_docnotes_10.0.html
Release Notes	<i>prod_os_relnotes_version.html/txt</i>	ids_unix_relnotes_10.0.txt
Machine Notes	<i>prod_machine_notes_version.txt</i>	ids_machine_notes_10.0.txt
Fixed and Known Defects File	<i>prod_defects_version.txt</i>	ids_defects_10.0.txt client_defects_2.90.txt
	<i>ids_win_fixed_and_known_defects_version.txt</i>	ids_win_fixed_and_known_defects_10.0.txt

Informix Error Messages

This file is a comprehensive index of error messages and their corrective actions for the Informix products and version numbers.

On UNIX platforms, use the **finderr** command to read the error messages and their corrective actions.

Dynamic Server

On Windows, use the Informix Error Messages utility to read error messages and their corrective actions. To display this utility, choose **Start > Programs > IBM Informix Dynamic Server *version* > Informix Error Messages** from the taskbar.

End of Dynamic Server

You can also access these files from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>.

Manuals

Online Manuals

A CD that contains your manuals in electronic format is provided with your IBM Informix products. You can install the documentation or access it directly from the CD. For information about how to install, read, and print online manuals, see the installation insert that accompanies your CD. You can also obtain the same online manuals from the IBM Informix Online Documentation site at <http://www.ibm.com/software/data/informix/pubs/library/>.

Printed Manuals

To order hardcopy manuals, contact your sales representative or visit the IBM Publications Center Web site at <http://www.ibm.com/software/howtobuy/data.html>.

Online Help

IBM Informix online help, provided with each graphical user interface (GUI), displays information about those interfaces and the functions that they perform. Use the help facilities that each GUI provides to display the online help.

Accessibility

IBM is committed to making our documentation accessible to persons with disabilities. Our books are available in HTML format so that they can be accessed with assistive technology such as screen reader software. The syntax diagrams in our manuals are available in dotted decimal format, which is an accessible format that is available only if you are using a screen reader. For more information about the dotted decimal format, see the Accessibility appendix.

IBM Informix Dynamic Server Version 10.0 and CSDK Version 2.90 Documentation Set

The following tables list the manuals that are part of the IBM Informix Dynamic Server, Version 10.0 and the CSDK Version 2.90, documentation set. PDF and HTML versions of these manuals are available at <http://www.ibm.com/software/data/informix/pubs/library/>. You can order hardcopy versions of these manuals from the IBM Publications Center at <http://www.ibm.com/software/howtobuy/data.html>.

Table 1. Database Server Manuals

Manual	Subject
Administrator's Guide	Understanding, configuring, and administering your database server.
Administrator's Reference	Reference material for Informix Dynamic Server, such as the syntax of database server utilities onmode and onstat , and descriptions of configuration parameters, the sysmasters tables, and logical-log records.
Backup and Restore Guide	The concepts and methods you need to understand when you use the ON-Bar and ontape utilities to back up and restore data.
DB-Access User's Guide	Using the DB-Access utility to access, modify, and retrieve data from Informix databases.
DataBlade API Function Reference	The DataBlade API functions and the subset of ESQL/C functions that the DataBlade API supports. You can use the DataBlade API to develop client LIBMI applications and C user-defined routines that access data in Informix databases.
DataBlade API Programmer's Guide	The DataBlade API, which is the C-language application-programming interface provided with Dynamic Server. You use the DataBlade API to develop client and server applications that access data stored in Informix databases.
Database Design and Implementation Guide	Designing, implementing, and managing your Informix databases.
Enterprise Replication Guide	How to design, implement, and manage an Enterprise Replication system to replicate data between multiple database servers.
Error Messages file	Causes and solutions for numbered error messages you might receive when you work with IBM Informix products.
Getting Started Guide	Describes the products bundled with IBM Informix Dynamic Server and interoperability with other IBM products. Summarizes important features of Dynamic Server and the new features for each version.
Guide to SQL: Reference	Information about Informix databases, data types, system catalog tables, environment variables, and the <code>stores_demo</code> demonstration database.
Guide to SQL: Syntax	Detailed descriptions of the syntax for all Informix SQL and SPL statements.
Guide to SQL: Tutorial	A tutorial on SQL, as implemented by Informix products, that describes the basic ideas and terms that are used when you work with a relational database.
High-Performance Loader User's Guide	Accessing and using the High-Performance Loader (HPL), to load and unload large quantities of data to and from Informix databases.
Installation Guide for Microsoft Windows	Instructions for installing IBM Informix Dynamic Server on Windows.
Installation Guide for UNIX and Linux	Instructions for installing IBM Informix Dynamic Server on UNIX and Linux.

Table 1. Database Server Manuals (continued)

Manual	Subject
J/Foundation Developer's Guide	Writing user-defined routines (UDRs) in the Java programming language for Informix Dynamic Server with J/Foundation.
Large Object Locator DataBlade Module User's Guide	Using the Large Object Locator, a foundation DataBlade module that can be used by other modules that create or store large-object data. The Large Object Locator enables you to create a single consistent interface to large objects and extends the concept of large objects to include data stored outside the database.
Migration Guide	Conversion to and reversion from the latest versions of Informix database servers. Migration between different Informix database servers.
Optical Subsystem Guide	The Optical Subsystem, a utility that supports the storage of BYTE and TEXT data on optical disk.
Performance Guide	Configuring and operating IBM Informix Dynamic Server to achieve optimum performance.
R-Tree Index User's Guide	Creating R-tree indexes on appropriate data types, creating new operator classes that use the R-tree access method, and managing databases that use the R-tree secondary access method.
SNMP Subagent Guide	The IBM Informix subagent that allows a Simple Network Management Protocol (SNMP) network manager to monitor the status of Informix servers.
Storage Manager Administrator's Guide	Informix Storage Manager (ISM), which manages storage devices and media for your Informix database server.
Trusted Facility Guide	The secure-auditing capabilities of Dynamic Server, including the creation and maintenance of audit logs.
User-Defined Routines and Data Types Developer's Guide	How to define new data types and enable user-defined routines (UDRs) to extend IBM Informix Dynamic Server.
Virtual-Index Interface Programmer's Guide	Creating a secondary access method (index) with the Virtual-Index Interface (VII) to extend the built-in indexing schemes of IBM Informix Dynamic Server. Typically used with a DataBlade module.
Virtual-Table Interface Programmer's Guide	Creating a primary access method with the Virtual-Table Interface (VTI) so that users have a single SQL interface to Informix tables and to data that does not conform to the storage scheme of Informix Dynamic Server.

Table 2. Client/Connectivity Manuals

Manual	Subject
Client Products Installation Guide	Installing IBM Informix Client Software Developer's Kit (Client SDK) and IBM Informix Connect on computers that use UNIX, Linux, and Windows.
Embedded SQLJ User's Guide	Using IBM Informix Embedded SQLJ to embed SQL statements in Java programs.

Table 2. Client/Connectivity Manuals (continued)

Manual	Subject
ESQL/C Programmer's Manual	The IBM Informix implementation of embedded SQL for C.
GLS User's Guide	The Global Language Support (GLS) feature, which allows IBM Informix APIs and database servers to handle different languages, cultural conventions, and code sets.
JDBC Driver Programmer's Guide	Installing and using Informix JDBC Driver to connect to an Informix database from within a Java application or applet.
.NET Provider Reference Guide	Using Informix .NET Provider to enable .NET client applications to access and manipulate data in Informix databases.
ODBC Driver Programmer's Manual	Using the Informix ODBC Driver API to access an Informix database and interact with the Informix database server.
OLE DB Provider Programmer's Guide	Installing and configuring Informix OLE DB Provider to enable client applications, such as ActiveX Data Object (ADO) applications and Web pages, to access data on an Informix server.
Object Interface for C++ Programmer's Guide	The architecture of the C++ object interface and a complete class reference.

Table 3. DataBlade Developer's Kit Manuals

Manual	Subject
DataBlade Developer's Kit User's Guide	Developing and packaging DataBlade modules using BladeSmith and BladePack.
DataBlade Module Development Overview	Basic orientation for developing DataBlade modules. Includes an example illustrating the development of a DataBlade module.
DataBlade Module Installation and Registration Guide	Installing DataBlade modules and using BladeManager to manage DataBlade modules in Informix databases.

Compliance with Industry Standards

The American National Standards Institute (ANSI) and the International Organization of Standardization (ISO) have jointly established a set of industry standards for the Structured Query Language (SQL). IBM Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992. In addition, many features of IBM Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL Common Applications Environment (CAE) standards.

IBM Welcomes Your Comments

We want to know about any corrections or clarifications that you would find useful in our manuals, which will help us improve future versions. Include the following information:

- The name and version of the manual that you are using
- Section and page number
- Your suggestions about the manual

Send your comments to us at the following email address:

`docinf@us.ibm.com`

This email address is reserved for reporting errors and omissions in our documentation. For immediate help with a technical problem, contact IBM Technical Support.

We appreciate your suggestions.

Chapter 1. Concepts

Features of Java User-Defined Routines	1-1
Java Virtual Processors	1-1
Thread Scheduling	1-2
Query Parallelization	1-2
System Catalog Tables	1-2

In This Chapter

This chapter introduces the infrastructure for creating and executing user-defined routines (UDRs) and applications that you write in Java to run in the Informix database server.

This chapter provides the following information:

- Basic characteristics of Java UDRs
- Basic architecture for executing Java UDRs in the database server
- Impact of Java UDRs on the database server system catalog tables

For general information on the purpose and the process of developing UDRs for the database server, refer to the *IBM Informix: User-Defined Routines and Data Types Developer's Guide*. For information on how to access databases from Java UDRs, refer to the *IBM Informix: JDBC Driver Programmer's Guide*.

Features of Java User-Defined Routines

The Informix database server provides the infrastructure to support Java UDRs. The database server binds SQL UDR signatures to Java executables and provides mapping between SQL data values and Java objects so that the database server can pass parameters and retrieve returned results.

The Informix database server also provides support for data type extensibility and sophisticated error handling.

Java Virtual Processors

Java UDRs execute on specialized virtual processors called *Java virtual processors* (JVPs). A Java Virtual Machine (JVM) is embedded in the code of each JVP.

The JVPs are responsible for executing all server-based Java UDRs and applications. Although the JVPs are mainly used for Java-related computation,

they have the same capabilities as a CPU VP, and they can process all types of SQL queries. This eliminates the need to ship Java-related queries back and forth between CPU VPs and JVPs.

Thread Scheduling

When the JVP starts the JVM, the entire database server component is thought of as running on one particular Java thread, called the *main thread*. The JVM controls the scheduling of Java threads and the database server scheduler multiplexes Informix threads on top of the Java main thread. In other words, the Informix thread package is stacked on top of the Java thread package.

Query Parallelization

While Java applications use threads for parallelism, the Informix database server uses threads for overlapping latency. That is, Informix threads run concurrently but not in parallel. To parallelize a query, the database server must spread the work among multiple virtual processors.

UNIX Only

Consequently, the database server must have multiple instances of JVPs to make parallel calls to UDRs written in Java code. Because the JVMs embedded in different VPs do not share states, you cannot store global states using Java class variables. All global states must be stored in the database to be consistent. The only guarantee from the database server is that any given UDR instance executes from start to finish on the same VP. The database server enforces a round-robin scheduling policy where the UDR instances are spread over the JVPs before they start executing.

End of UNIX Only

Windows 2000 Only

The consistency of multiple JVMs is not an issue on the Windows platform because all VPs are mapped to kernel threads instead of processes. Because all VPs share the same process space, you do not need to start multiple instances of the JVM.

End of Windows 2000 Only

System Catalog Tables

The `sysroutinelangs`, `syslangauth`, and `sysprocedures` system catalog tables contain information about the UDRs written in Java code.

The **sysroutinelangs** table lists the programming languages that you can use to write UDRs. The table gives the names of the language initialization functions and the path for the language library.

The **syslangauth** table specifies who is allowed to use the language. For Java code, the default is the database administrator. For information about how to modify the use privileges, refer to the GRANT statement in the *IBM Informix: Guide to SQL Syntax*.

The **sysprocedures** table gives information about both built-in routines and routines that you define.

For more information about these system catalog tables, refer to “Finding Information about User-Defined Routines” on page 4-18 and to the *IBM Informix: Guide to SQL Reference*.

Chapter 2. Preparing for Java Support

Installing the JDBC Driver	2-1
Configuring Java Support	2-2
Creating an sbospace	2-2
Creating the JVP Properties File	2-2
Setting Configuration Parameters	2-3
Setting Environment Variables	2-4
GLS Support	2-5
NEWLOCALE and NEWCODESET Connection Properties	2-5
DBCENTURY Environment Variable	2-6

In This Chapter

This chapter describes how to install and configure the database server to provide UDRs written in Java code. To create and use UDRs written in Java code, you must install the following software:

- IBM Informix Dynamic Server with J/Foundation
- The Java Development Kit (JDK), Version 1.4.x

If you do not plan to develop Java UDRs, you do not need to install the JDK. J/Foundation includes a tested version of the *Java Runtime Environment* (JRE) to execute Java UDRs. You need to install the JDK only if you need to compile new Java source code. For more information on where to obtain the JDK, see the machine notes for your platform.

You might also want to install the DataBlade Developers Kit (DBDK), Version 4.0 or greater, to facilitate development of UDRs in Java code.

For more detailed information on the required software, refer to the release notes.

Installing the JDBC Driver

J/Foundation includes the IBM Informix JDBC Driver. The IBM Informix JDBC Driver contains Java classes and shared-object files that allow you to write UDRs in Java code. The installation procedure installs these binaries in **\$INFORMIXDIR/extend/krakatoa**.

For more information, refer to the machine notes file.

Configuring Java Support

The basic configuration procedure for an Informix database server is covered in the *IBM Informix: Dynamic Server Administrator's Guide*. Configuring the database server to support Java code requires several additional steps. You might find it convenient to configure the database server without Java code and then modify it to add Java support.

Preparing to use Java code with the database server requires these additions to the basic configuration procedure:

- Create an sbspace to hold the Java JAR files.
- Create the JVP properties file.
- Add (or modify) the Java configuration parameters in the ONCONFIG configuration file.
- Set environment variables.

`$INFORMIXDIR/extend/krakatoa` is your *jvphome*. You need to include this path in several places as you prepare J/Foundation.

Creating an sbspace

The database server stores Java JAR files as smart large objects in the system default sbspace. If you do not already have a default sbspace, you must create one. For example, the following command creates an sbspace called **mysbspace**:

```
onspaces -c -S mysbspace -g 5 -p /dev/raw_dev1 -o 500 -s 20000 -m /dev/raw_dev2 500
```

For information about the **onspaces** command, refer to the *IBM Informix: Administrator's Reference*.

After you create the sbspace, set the SBSPACENAME configuration parameter in the ONCONFIG file to the name that you gave to the sbspace (**mysbspace** in the preceding example).

JAR files coexist in the system default sbspace with other smart large objects that you store in that space. When you choose the size for your default sbspace, you need to consider how much space those objects require, as well as the number and size of the JAR files that you plan to install.

Creating the JVP Properties File

A *JVP properties file* contains property settings that control various runtime behaviors of the Java virtual processor. The JVPPROPFIL configuration parameter specifies the path to the properties file. When you initialize the database server, the JVP initializes the environment based on the settings in the JVP property file. The **.jvpprops.template** file in the `$INFORMIXDIR/extend/krakatoa` directory documents the properties that you can set.

To prepare the JVP properties file:

1. Copy the JVP properties template file, `jvphome/.jvpprops.template` into `jvphome/.jvpprops` where `jvphome` is the directory `$INFORMIXDIR/extend/krakatoa`.
2. Edit `.jvpprops` to change the trace level or other properties if necessary.
3. Set the `JVPPROFILE` configuration parameter to `jvphome/.jvpprops`.

A sample properties file might contain the following items:

```
JVP.trace.settings:JVP=2
JVP.trace.verbose:1
JVP.trace.timestampformat:HH:MM
JVP.splitLog:1000
JVP.monitor.port: 10000
```

The database server provides a fixed set of system trace events such as UDR sequence initialization, activation, and shutdown. You can also generate application-specific traces. For more information, see the description of the `UDRTraceable` class in “The `com.informix.udr.UDRTraceable`” on page 4-7.

Setting Configuration Parameters

The `ONCONFIG` configuration file (`$INFORMIXDIR/etc/$ONCONFIG`) includes the following configuration parameters that affect Java code:

- `JDKVERSION`
- `JVPPROFILE`
- `JVMTHREAD`
- `JVPCLASSPATH`
- `JVPHOME`
- `JVPJAVALIB`
- `JVPJAVAVM`
- `JVPLOGFILE`
- `VPCLASS`

UNIX Only

The following example shows sample settings for the Java-related configuration parameters on a UNIX Solaris system. In this example, `jvphome` is `$INFORMIXDIR/extend/krakatoa`.

```
JVPHOME      jvphome
JVPLOGFILE   jvphome/jvp.log
JVPPROFILE   jvphome/.jvpprops
JVPJAVAVM    java_g:net_g:zip_g:mmedia_g:jpeg_g:
              sysresource_g:agent_g
VPCLASS      jvp,num=1
JDKVERSION   1.4
```

```
JVMTHREAD      native
JVPJAVALIB    /lib/sparc/native_threads
JVPCLASSPATH  jvphome/krakatoa_g.jar:jvphome/jdbc_g.jar
```

In this example, JVPJAVAVM and JVPCLASSPATH are set appropriately for debug mode. To run in nondebug mode, remove all the `_g` suffixes.

For more information, refer to Chapter 3, “Configuration Parameters,” on page 3-1. For information about specific configuration parameter settings on your platform, refer to `$INFORMIXDIR/etc/onconfig.std`.

End of UNIX Only

Setting Environment Variables

You do not need any extra environment variables to execute UDRs written in Java code. However, if you are developing Java UDRs, you must include `jvphome/krakatoa.jar` in your `CLASSPATH` environment variable so that JDK can compile the Java source files that use Informix Java packages.

The following sections describe the runtime environment variables that you can set.

AFDEBUG

Set `AFDEBUG` to 1 to create files to hold verbose garbage collection messages from the JVM. You can also set the configuration parameter `AFCRASH` to `0x00000010` to achieve the same result.

JVM_MAX_HEAP_SIZE

Set the environment variable `JVM_MAX_HEAP_SIZE` to configure the heap size for the JVM. The default heap size is 16 megabytes. You can set this variable to the maximum heap size needed for the JVM, depending on the estimated requirements of the application.

JAR_TEMP_PATH

Set the `JAR_TEMP_PATH` environment variable to specify a local file-system location where jar management procedures such as `install_jar` and `replace_jar` can store JAR files temporarily. This directory must have read and write permissions for the user who brings up the database server. If the `JAR_TEMP_PATH` environment variable is not set, temporary copies of JAR files are stored in the `/tmp` directory of the local file system for the database server.

JAVA_COMPILER

To turn off just-in-time (JIT) compilation, set the `JAVA_COMPILER` environment variable to `NONE` or `none`. For more information on JIT compilation, see the Java documentation from Sun Microsystems.

GLS Support

When the database server starts a UDR, the routine runs in the locale that `DB_LOCALE` specifies. Consequently, the database server automatically converts parameters, return values, and output values between the `DB_LOCALE` code set and the Unicode code set so that Java code can use the values.

However, when a Java UDR creates a JDBC connection to the database server for access through SQL, you can set `DB_LOCALE` into the connection URL to control conversions and formatting between the Unicode code set and the code set of the database server locale. This setting of `DB_LOCALE` overrides any environment settings. In fact, `DB_LOCALE` does not need to be set in the environment. Similarly, you can also set `DBDATE`, `GL_DATE`, and `DBCENTURY` into the URL connection to control date conversion and formatting.

For example, when a UDR sends string or date data to the database server in an insert, the database server converts the data from Unicode to the locale that `DB_LOCALE` specifies, or it interprets dates and intervals using your `DBDATE` or `GL_DATE` setting.

When the database server returns data to the Java UDR, the database server does the opposite conversion, so Java code sees only Unicode.

NEWLOCALE and NEWCODESET Connection Properties

IBM Informix JDBC Driver uses the JDK internationalization API to manipulate international data. The classes and methods in this API take a JDK locale or encoding as a parameter. Because the Informix `DB_LOCALE` and `CLIENT_LOCALE` properties specify the locale and code set based on Informix names, these Informix names are mapped to the JDK names. For example, the Informix and JDK names for the ASCII code set are `8859-1` and `8859_1` respectively. IBM Informix JDBC Driver internally maps `8859-1` to `8859_1` and uses the appropriate JDK name in the JDK classes and methods.

Two new connection properties, `NEWLOCALE` and `NEWCODESET`, enable you to specify a locale or code set that is not yet mapped in the internal tables of the JDBC driver.

The `NEWLOCALE` and `NEWCODESET` properties have the following formats:

```
NEWLOCALE=<JDK locale>,<Ifx locale>:<JDK locale>,<Ifx locale>...
NEWCODESET=<JDK encoding>,<Ifx codeset name>,<Ifx codeset number>:<JDK encoding>,<Ifx codeset name>,<Ifx codeset number>...
```

The following example shows a URL that uses these properties. (You must specify a valid URL on a single line.)

```
jdbc:informix-sqli://myhost:1533:informixserver=myserver;user=myname;  
password=mpasswd;NEWLOCALE=en_us,en_us;NEWCODESET=8859_1,8859-1,819;
```

There is no limit to the number of locale or code-set mappings that you can specify. If you specify an incorrect number of parameters or values, you get a message that says, “Locale Not Supported” or “Encoding or Code Set Not Supported.” If you set these properties in the URL or in an **IfmxDataSource** object, the new values in NEWLOCALE and NEWCODESET override the values in the JDBC internal tables. For example, if JDBC already maps 8859-1 to 8859_1, but you specify NEWCODESET=8888,8859-1,819, the new value, 8888, is used for the code-set conversion.

DBCENTURY Environment Variable

If a String represents a DATE or a DATETIME value that has less than a three-digit year value, the IBM Informix JDBC Driver uses the **DBCENTURY** environment variable to determine the correct four-digit year and performs a String-to-DATE or -DATETIME conversion. For the **DBCENTURY** settings, the algorithms used, and examples, see the section on environment variables in the *IBM Informix: Guide to SQL Reference*.

The following table summarizes the affected methods and the conditions under which they are affected.

Method	Condition
IfxPreparedStatement.setString(String)	The target column is SQLDATE or SQLDTIME.
IfxPreparedStatement.setObject(String)	The target column is SQLDATE or SQLDTIME.
IfxPreparedStatement.IfxSetObject(String)	The target column is SQLDATE or SQLDTIME.
IfxResultSet.getDate()	The source column is a String type.
IfxResultSet.getTimestamp	The source column is a String type.
IfxResultSet.updateString(String)	The target column is SQLDATE or SQLDTIME.
IfxResultSet.updateObject(String)	The target column is SQLDATE or SQLDTIME.

The following example shows a URL that uses the **DBCENTURY** environment variable:

```
jdbc:informix-sqli://myhost:1533:informixserver=myserver;user=myname;password=
mypasswd;DBCENTURY=F;
```

You must specify a valid URL on a single line.

Chapter 3. Configuration Parameters

JVPARGS	3-1
JDKVERSION	3-2
JVPCLASSPATH	3-2
JVPHOME	3-3
JVPJAVAHOME	3-3
JVPJAVALIB	3-4
JVPJAVAVM	3-4
JVPLOGFILE	3-5
JVPPROFILE	3-6
SBSACENAME	3-6
VPCCLASS JVP	3-6

In This Chapter

This chapter documents the configuration parameters that you need to set to use UDRs written in Java code. Set these parameters in the database server configuration file (the ONCONFIG file).

For a sample environment that configuration parameters establish, see the release notes file.

JVPARGS

onconfig.std

value None

takes effect When shared memory is initialized

The JVPARGS configuration parameter provides an easy way for you to set Java VM options.

Use a semicolon to separate options. For example, if you want to change `Xms` and `Xmx` to 32m, you can set those options with the JVPARGS parameter, as the following example shows:

```
JVPARGS -Xms32m;-Xmx32m
```

If you want to see gc information to determine whether you need to increase `ms` or `mx`, you can set JVPARGS, as the following example shows:

```
JVPARGS -verbose:gc
```

For more information on Java VM options, refer to your Java documentation.

JDKVERSION

onconfig.std

value 1.4

range of values For this release, the valid values are 1.4, 1.3, and 1.2.

takes effect When shared memory is initialized

JDKVERSION is the major version of the JDK or JRE release. That is, the version number does not include *x* when the version is JDK 1.4.*x*.

This parameter is required if the number of JVPs (set in VPCLASS JVP) is greater than 0.

JVPCLASSPATH

onconfig.std

value **/usr/informix/extend/krakatoa/krakatoa_g.jar:
/usr/informix/extend/krakatoa/jdbc_g.jar**

takes effect When shared memory is initialized

The JVPCLASSPATH configuration parameter is the initial Java classpath setting. You must modify the default setting in the configuration file by replacing **/usr/informix/extend/krakatoa** with **JVPHOME_path**, the pathname in your JVPHOME configuration parameter.

JVPHOME_path/krakatoa_g.jar:JVPHOME_path/jdbc_g.jar

If you do not require the debug versions of the jar files, use the following JVPCLASSPATH setting:

JVPHOME_path/krakatoa.jar:JVPHOME_path/jdbc.jar

The total number of characters available for specifying configuration values in the ONCONFIG file is 256. The database server imposes this limit.

To specify more than 256 characters for the value of the JVPCLASSPATH parameter, you can store the value in a file and specify the keyword **file:** on the parameter, followed by the filename. For example, if you set the path in a file called **classpath_fl** in the directory **/u/informix/iif2000/extend/java**, you can specify the JVPCLASSPATH parameter, as the following example shows:

JVPCLASSPATH file:/u/informix/iif2000/extend/java/classpath_fl

You must specify the complete value for JVPCLASSPATH on one line in the file, just as you would normally on the configuration parameter. Do not

include the parameter name JVPCLASSPATH again. The database server considers the first carriage return in the line to be the terminating carriage return for the pathname.

The JVPCLASSPATH parameter is required if the number of JVPs (set in VPCLASS JVP parameter) is greater than 0.

Tip: The JVP ignores the **CLASSPATH** environment variable. However, you must set the **CLASSPATH** environment variable so that you can compile your UDRs.

JVPHOME

onconfig.std

value **/usr/informix/extend/krakatoa**

takes effect When shared memory is initialized

The JVPHOME configuration parameter specifies the directory where the classes of the IBM Informix JDBC Driver are installed. To modify the default setting in the configuration file, replace **/usr/informix** with the pathname of your **\$INFORMIXDIR**.

The JVPHOME value, *JVPHOME_path*, is used in several configuration parameters. If the JVPHOME location changes, you must change the configuration settings of all parameters that use the JVPHOME value.

This parameter is required if the number of JVPs (set in VPCLASS JVP) is greater than 0.

JVPJAVAHOME

onconfig.std

value **/usr/informix/extend/krakatoa/jre/**

takes effect When shared memory is initialized

The JVPJAVAHOME configuration parameter specifies the directory where the JRE for the database server is installed. The database server includes a tested version of the JRE. The default location for the JRE is in **/usr/informix/extend/krakatoa/jre/**. To modify the default setting in the configuration file, replace **/usr/informix/extend/krakatoa/jre** with the pathname setting of **\$INFORMIXDIR**, followed by **/extend/krakatoa/jre**.

This parameter is required if the number of JVPs (set in VPCLASS JVP) is greater than 0.

If you want to use a stand-alone JVM, without a JVP, install the JDK on your platform and use the JVM that is included.

JVPJAVALIB

onconfig.std

value *platform-specific value*

takes effect When shared memory is initialized

The JVPJAVALIB configuration parameter specifies the path from **\$JVPJAVAHOME** to the location of the JVM libraries.

The value of this parameter is platform dependent. To find the proper value for JVPJAVALIB, refer to the machine and release notes.

This parameter is required if the number of JVPs (set in VPCLASS JVP) is greater than 0.

JVPJAVAVM

onconfig.std

value *platform-specific value*

separators colon (UNIX) and semicolon (Windows)

takes effect When shared memory is initialized

The JVPJAVAVM configuration parameter lists the JVM libraries that the database server should load. The names in this list exclude the **lib** prefix and **.so** or **.dll** suffix. Entries in the list are separated by colons.

This parameter is required if the number of JVPs (set in VPCLASS JVP) is greater than 0.

UNIX Only

For example, for UNIX Solaris, use the following value for JVPJAVAVM if you are using a debug version of the JDBC driver:

```
hpi_g:server_g:verify_g:java_g:net_g:jpeg_g
```

End of UNIX Only

If you use a nondebug JDBC driver, you can use the nondebug JDK libraries for better performance. Set JVPJAVAVM to the following:

```
hpi:server:verify:javag:net:jpeg
```

Windows 2000 Only

For Windows, use a semicolon to separate values. Use the following value for VPJAVAVM if you are using a debug version of the JDBC driver:

```
hpi_g;server_g;verify_g;java_g;net_g;jpeg_g
```

If you use a nondebug JDBC driver, you can use the nondebug JDK libraries for better performance. Set JVPJAVAVM to the following:

```
hpi;server;verify;javag;net;jpeg
```

End of Windows 2000 Only

The value of JVPJAVAVM is platform dependent. To find the proper value for JVPJAVAVM, refer to the machine and release notes.

JVPLOGFILE

onconfig.std

value **/usr/informix/jvp.log**

range of values Any valid complete filename

takes effect When shared memory is initialized

The database server can generate Java trace outputs and stack dumps. The database server writes this output to the Java VP log file.

The JVPLOGFILE configuration parameter specifies the path to the Java VP log file. This parameter is optional.

To change the location of the log file, change the value of the JVPLOGFILE configuration parameter. For example, the following parameter value sets the log file to **/u/sam/jvp.log**:

```
JVPLOGFILE /u/sam/jvp.log
```

If you do not specify a value for this parameter, the default value is derived from the **onconfig.std** file. If the JVPLOGFILE parameter is not present in the ONCONFIG file, the default file location is:

```
./jvp.log
```

where '.' is the current directory of the user who runs **oninit**.

JVPPROFILE

onconfig.std

value **/usr/informix/extend/krakatoa/.jvpprops**

takes effect When shared memory is initialized

The JVPPROFILE configuration parameter specifies the path to the Java VP properties file, if any. Set this parameter as follows, where *JVPHOME_path* is the value in your JVPHOME configuration parameter:

JVPHOME_path/.jvpprops

This parameter is optional.

SBSPACENAME

onconfig.std

value *blank*

takes effect When shared memory is initialized

refer to “Creating an sbspace” on page 2-2

The SBSPACENAME configuration parameter specifies the name of the system default sbspace. You must provide an sbspace where the database server can store the Java jar files.

This parameter is not exclusively for Java code. If your database tables include smart-large-object columns that do not explicitly specify a storage space, that data is stored in the sbspace that SBSPACENAME specifies.

For information about how to specify a storage space for smart large objects, refer to the CREATE TABLE statement in the *IBM Informix: Guide to SQL Syntax*. For more information about SBSPACENAME, refer to the *IBM Informix: Administrator's Reference*.

Tip: When you use UDRs written in Java code, create separate sbspaces for storing your smart large objects.

VPCLASS JVP

onconfig.std

value *not set*

range of values 0 and positive integers

takes effect When shared memory is initialized

The VPCLASS configuration parameter specifies the number of virtual processors to initialize for a given virtual-processor class. The JVP option of VPCLASS specifies the number of Java virtual processors that the database server should start.

This parameter is required to execute Java UDRs.

Set this option as follows, where *number* is the number of Java virtual processors:

```
VPCLASS JVP,num=number
```

The default value of *range* is 1. If you set the number of JVPs to zero (0), or if there is no VPCLASS parameter for the JVP class, execution of Java UDRs is disabled.

If you have not correctly installed and configured the software for Java in the server, the JVP fails to start when you start the database server. However, the database server itself continues to initialize normally. The main database log file contains a message that indicates the cause of the JVP failure.

For more information about the VPCLASS configuration parameter, refer to the *IBM Informix: Administrator's Reference*.

Chapter 4. Creating Java User-Defined Routines

Java User-Defined Routines	4-2
Creating a Java User-Defined Routine	4-3
Writing a Java User-Defined Routine	4-4
The com.informix.udr Package	4-5
The com.informix.udr.UDRManager	4-5
The com.informix.udr.UDREnv	4-5
The com.informix.udr.UDRLog	4-7
The com.informix.udr.UDRTraceable	4-7
Creating UDT-to-Java Mappings	4-8
Registering Java User-Defined Routines	4-9
Specifying the JVP	4-9
Using Routine Modifiers	4-10
Specifying the External Name	4-11
Using a Deployment Descriptor	4-13
Using a Manifest File	4-14
Compiling the Java Code	4-14
Installing a JAR File	4-15
Updating JAR Filenames	4-16
Executing the User-Defined Routine	4-16
Debugging and Tracing	4-16
Generating Log Messages.	4-16
Using the Administrative Tool	4-17
The threads vpid Option	4-18
The memory vpid Option.	4-18
Debugging a Java User-Defined Routine.	4-18
Traceable Events.	4-18
Finding Information about User-Defined Routines	4-18
Complying with SQLJ	4-19
Unsupported Modifiers	4-20
Unsupported Optional Modifiers	4-21

In This Chapter

A *user-defined routine* (UDR) is a routine that an SQL statement or another UDR can invoke. UDRs written in Java code use the server-side implementation of the IBM Informix JDBC Driver to communicate with the database server.

This chapter provides the following information about UDRs written in Java code:

- What tasks a UDR can perform
- How to create a UDR

Java User-Defined Routines

The behaviors of installing and invoking UDRs written in Java code follow the SQLJ: SQL Routines specification. Every UDR written in Java code maps to an external Java static method whose class resides in a Java Archive (JAR) file that was installed in a database. The SQL-to-Java data type mapping is done according to the JDBC specification.

UDRs can be user-defined functions or user-defined procedures, which can return values or not, as follows:

- A *user-defined function* returns one or more values and therefore can be used in SQL expressions.

For example, the following query returns the results of a UDR called **area()** as part of the query results:

```
SELECT diameter, area(diameter) FROM shapes
WHERE diameter > 6
```

- A *user-defined procedure* is a routine that optionally accepts a set of arguments and does not return any values.

A procedure *cannot* be used in SQL expressions because it does not return a value. However, you can call it directly, as the following example shows:

```
EXECUTE PROCEDURE myproc(1, 5)
```

You can also call user-defined procedures within triggers.

For general information about UDRs, refer to the *IBM Informix: User-Defined Routines and Data Types Developer's Guide*.

UDRs written in Java code can perform the following tasks.

Type of UDR	Purpose
End-user routine	A UDR that performs some common task for an end user
User-defined aggregate	A UDR that calculates an aggregate value on a PROCEDURE particular column or value
Parallelizable UDR	A UDR that can run in parallel when executed within an SQL statement (UDRs that open JDBC connections cannot run in parallel.)
Cast function	A UDR that converts or casts one data type to another
Operator function	A UDR that implements some operator symbol (such as +, -, or /)

Iterator function	A user-defined function that returns more than one row of data Iterator functions written in Java code are supported using some Informix extensions.
Functional index	A UDR on which an index can be built
Opaque data type support function	A user-defined function that tells the database server how to handle the data of an opaque data type
Negator function	A function that calculates the <i>not</i> operation for a particular operator or function

You *cannot* use UDRs written in Java code for any of the following features:

- Commutator functions
- Cost functions
- Operator-class functions
- Selectivity functions
- User-defined statistics functions

Creating a Java User-Defined Routine

When you create a Java UDR, you need to write and compile the source code and then install the finished code in the database server.

To create a Java UDR:

1. Write the UDR, which can use the JDBC methods to interact with the database server.
2. If the UDR uses any user-defined data types (UDTs), for each UDT write a Java class that translates between the database server and Java representation of the type.
This class should implement the **SQLData** interface. For information about **SQLData**, refer to the JDBC 2.0 specification.
3. Write the CREATE FUNCTION or CREATE PROCEDURE statement for registering the UDR.
4. Write the deployment descriptor, which contains the SQL statements for registering the UDR.
5. Prepare the manifest file.
6. Compile the Java source files and collect the compiled code into a JAR file.
7. Create a JAR file that contains the classes, deployment descriptor, and manifest file.

8. Install the JAR file that contains the UDR in the current database.
9. Execute the UDR.
10. Use tracing and the debugging features to work out any problems in the UDR.
11. Optimize performance of the UDR.

For general information on how to develop a UDR, refer to the *IBM Informix: User-Defined Routines and Data Types Developer's Guide*. The following sections briefly describe each of these steps in the development of a UDR.

Tip: It is recommended that you use the *DataBlade Developers Kit (DBDK)*, Version 4.0 or later, to help write UDRs in Java code. DBDK enforces standards that facilitate migration between different versions of the database server.

Writing a Java User-Defined Routine

Java UDRs can use the following packages, interfaces, classes, and methods:

- Java packages
UDRs can use all the basic nongraphic Java packages that are in the JDK. That is, UDRs can use **java.util.***, **java.io.***, **java.net.***, **java.rmi.***, and so on. UDRs cannot use **java.awt.***, **java.applet.*** and other user-interface packages. For more information on these packages, see the JDK documentation.
- Java Database Connectivity (JDBC) 1.0 API
UDRs can use the JDBC 1.0 API to access the database. For more information, see “JDBC 1.0 API” on page 5-3.
The **\$INFORMIXDIR/extend/krakatoa/examples.tar** file of online examples includes a sample of JDBC in a UDR in **JDBC.java**.
- Informix JDBC extensions
UDRs can also use Informix extensions to JDBC 1.0 to access some JDBC 2.0 functionality. For more information, see Chapter 5, “The IBM Informix JDBC Driver,” on page 5-1.
- Informix extensions for UDRs written in Java code
Certain Informix extensions are available to applications that need to exploit the capabilities of the database server. The Informix extensions reside in the **com.informix.udr package**.

The Informix **com.informix.udr** package provides extensions to SQLJ that allow applications to exploit the capabilities of Dynamic Server. Such extensions include logging, tracing, iterator support, and invocation-state management.

The `com.informix.udr` Package

The `com.informix.udr` package contains the following public interfaces:

- The `com.informix.udr.UDRManager`
- The `com.informix.udr.UDREnv`
- The `com.informix.udr.UDRLog`
- The `com.informix.udr.UDRTraceable`

The following sections describe each of these Informix-specific extensions.

The `com.informix.udr.UDRManager`

The `UDRManager` class provides a method for a UDR instance to obtain its `UDREnv` object. This class is defined as follows:

```
public class UDRManager
{
    static UDREnv getUDREnv();
}
```

The SQLJ: SQL Routines specification, which describes how to use static Java methods as database UDRs, does not provide a mechanism to save the user state across invocations. The `UDREnv` interface is a provided interface that maintains state information. You can use this state information, for example, to write iterator UDRs. The `UDREnv` object is maintained by the thread that manages the execution of the static method that represents the UDR. Therefore, if the UDR forks its own threads, the `UDRManager.getUDREnv` method cannot be directly used by those secondary threads of the UDR. The UDR must explicitly pass the `UDREnv` object to the secondary threads that it creates.

The `com.informix.udr.UDREnv`

The `UDREnv` interface consists of methods for accessing and manipulating the routine state of the UDR. It exposes a subset of the routine-state information in the `MI_FPARAM` structure (which holds routine-state information for C UDRs). It also contains some utilities related to the JVP, such as logging and tracing.

The online examples in `$INFORMIXDIR/extend/krakatoa/examples.tar` include an example of the `UDREnv` class in `Env.java`.

The `UDREnv` interface is defined as follows:

```
public interface UDREnv
{
    // Information about the UDR signature

    String getName();
    String[] getParamTypeName();
    String getReturnTypeName();
}
```

```

// For maintaining state across UDR invocations

void setUDRState (Object state);
Object getUDRState();

// For set/iterator processing

public static final int UDR_SET_INIT = 1;
public static final int UDR_SET_RETONE = 2;
public static final int UDR_SET_END = 3;
int getSetIterationState();
void setSetIterationIsDone(boolean value);

// Logging and Tracing

UDRTraceable getTraceable();
UDRLog getLog();
}

```

The **getName()** method returns the name of the UDR as it is registered in the database.

The **getParamTypeName()** and **getReturnTypeName()** methods return the SQL data type names for the UDR arguments and the return value, respectively.

If you are using JDBC2.0, use the **getUDRs()** method of the **java.sql.DatabaseMetaData** class to obtain more information about a data type.

The **setUDRState()** method sets the user-state pointer for the UDR. It stores a given object in the context of the UDR instance. The object might contain states that are shared across UDR invocations (such as a JDBC connection handle or a **UDRLog** object). The **getUDRState()** method returns the object set by the latest call to **setUDRState()**.

The **getSetIterationState()** method retrieves the iterator status for an iterator function. (This method is analogous to the C-language accessor **mi_fp_request** for set iterators.) This method returns one of the following values.

Iterator-Status Constant	Meaning	Use
UDR_SET_INIT	This is the <i>first</i> time that the iterator function is called.	Initialize the user state for the iterator function.
UDR_SET_RETONE	This is an actual iteration of the iterator function.	Return items of the active set, one per iteration.
UDR_SET_END	This is the last time that the iterator function is called.	Free any resources associated with the user state.

The **setSetIterationIsDone()** method sets the iterator-completion flag for an iterator function. Use the **setSetIterationIsDone()** method to tell the database server whether the current iterator function has reached its end condition. An *end condition* indicates that the generation of the active set is complete. The database server calls the iterator function with the UDR_SET_RETONE iterator-status value as long as the end condition has *not* been set.

The **getLog()** method returns a **UDRLog** interface for logging uses. For more information on the **UDRLog** interface, see “The com.informix.udr.UDRLog” on page 4-7.

The **getTraceable()** method returns a **UDRTraceable** interface for the UDRs to use. For more information on the **UDRTraceable** interface, see “The com.informix.udr.UDRTraceable” on page 4-7.

The com.informix.udr.UDRLog

The **UDRLog** interface provides a simple logging facility for a UDR. The **UDRLog** interface is defined as follows:

```
public interface UDRLog
{
    void log(String msg);
}
```

The interface defines a single method, **log()**, which takes a String argument and appends it to the JVP log file, which the JVPLOGFILE configuration parameter specifies. For more information, see “Generating Log Messages” on page 4-16.

The com.informix.udr.UDRTraceable

The **UDRTraceable** interface supports *zone-based* tracing. A trace zone is a conceptual code component. For example, you can put all UDRs in the same zone and all general-purpose Java applications in another. Each zone can have its own *trace level* that dictates the granularity of tracing. The zones form a hierarchy where subzones inherit the trace levels of their parents. You can define the zones, their hierarchical relationships, and trace levels with the following features:

- The settings in the JVP property file (which the JVPPROFILE configuration parameter specifies)
- Calls to the **UDRTraceable** methods at program execution time

The **UDRTraceable** interface is defined as follows:

```
public interface UDRTraceable extends Traceable
{
    public static final int TRACE_OFF = 0;
    public static final int TRACE_MINIMAL = 1;
    public static final int TRACE_COARSE = 2;
    public static final int TRACE_MEDIUM = 3;
}
```

```

public static final int TRACE_FINE = 4;
public static final int TRACE_SUPERFINE = 5;

int traceLevel(String zone);
void traceSet(String zone, int level);
void tracePrint(String zone, int level, String message);

```

The **traceLevel()** method returns the current trace-level setting for the given trace zone. The predefined trace levels are as follows.

Trace-Level Constant	Description
TRACE_OFF	No trace output is generated
TRACE_MINIMAL	Basic tracing
TRACE_COARSE	Coarse-grained tracing
TRACE_MEDIUM	Medium-grained tracing
TRACE_FINE	Fine-grained tracing
TRACE_SUPERFINE	For the trace sessions that require all possible details

The **traceSet()** method sets the specified trace zone to the specified trace level.

The **tracePrint()** method sends the specified message to the JVP log file if the trace zone has a trace level that is greater than or equal to the *level* parameter. The JVPLOGFILE configuration parameter specifies the JVP log-file name. For more information, see “Generating Log Messages” on page 4-16.

Creating UDT-to-Java Mappings

The routine manager needs a mapping between SQL data values and Java objects to be able to pass parameters to and retrieve return results from a UDR. The SQL-to-Java data type mapping is performed according to the JDBC specification. For built-in SQL data types, the routine manager can use mappings to existing JDBC data types.

For any UDTs that your UDR uses, you must create mappings. You can use the following UDTs in UDRs written in Java code.

User-Defined Data Type	SQL Statement
Distinct data type	CREATE DISTINCT TYPE
Opaque data type	CREATE OPAQUE TYPE

Warning: You *cannot* use row or collection data types in UDRs written in Java code.

To create the mapping between a user-defined SQL data type and a Java object:

1. Create a user-defined class that implements the **SQLData** interface. (For more information, refer to the JDBC 2.0 specification.)
2. Bind this user-defined class to the user-defined SQL data type using the **setUDTextName** built-in procedure.

Because the SQL statements that create UDTs do not currently provide a clause for specifying the external name of a UDT, you must define this mapping. Use the following built-in procedures with the EXECUTE PROCEDURE statement to define the mapping:

- **sqlj.setUDTextName()**

This procedure defines the mapping between a UDT and a Java data type.

- **sqlj.unsetUDTextName()**

This procedure removes the SQL-to-Java mapping and removes any cached copy of the Java class from database server shared memory.

For example:

```
-- Creating or removing UDT-to-Java Mappings
EXECUTE PROCEDURE sqlj.setUDTextName('udt_name',
    'class_name.udtname');
EXECUTE PROCEDURE sqlj.unsetUDTextName('udt_name');
```

The online examples in `$INFORMIXDIR/extend/krakatoa/examples.tar` include a sample implementation of a UDT written in Java code, **Circle.java**.

Registering Java User-Defined Routines

For a UDR to be invoked in an SQL statement, it must be registered in the current database. Use the CREATE FUNCTION and CREATE PROCEDURE statements to register UDRs. For details about SQLJ compliance, refer to “Complying with SQLJ” on page 4-19.

Tip: Place your SQL statements for registering UDRs written in Java code in a deployment descriptor file.

The following sections describe the Java-specific syntax of the CREATE FUNCTION and CREATE PROCEDURE statements that affects UDR registration. For information on the complete syntax of these SQL statements, see the *IBM Informix: Guide to SQL Syntax*.

Specifying the JVP

To execute, a UDR written in Java code must run in a JVP. The JVP is a predefined virtual-processor class that contains a JVM to interpret Java byte codes. Use the following syntax to specify that a UDR should execute in the JVP class:

```
WITH (class='jvp')
```

By default, most UDRs run in the CPU VP, which does *not* contain a JVM. However, a UDR written in Java code runs on a JVP by default. Therefore, the CLASS routine modifier is optional when you register a UDR written in Java code. To improve readability of your SQL statements, include the CLASS routine modifier when you register a UDR.

For example:

```
-- Specifying the JVP
CREATE PROCEDURE showusers()
  WITH (class='jvp')
  EXTERNAL NAME 'thisjar:admin.showusers()'
  LANGUAGE java;
```

Using Routine Modifiers

The routine modifiers that you specify in the WITH clause of the CREATE FUNCTION or CREATE PROCEDURE statement tell the database server about attributes of the UDR. The database server supports the following routine modifiers for UDRs.

Routine Modifier	Type of UDR
CLASS	Accesses to the JVP
HANDLESNULLS	Handles SQL null values as arguments
ITERATORS	Iterator function
NEGATOR	Negator function
NOT VARIANT	Might return cached results
PARALLELIZABLE	Parallelizable UDR
VARIANT	Returns different results when invoked with the same arguments

The following routine modifiers are C-language specific and do not apply to UDRs in Java code:

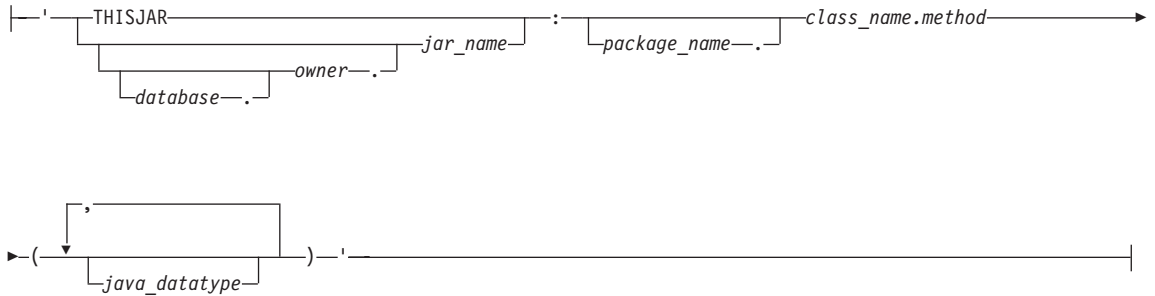
- COSTFUNC
- INTERNAL
- SELFUNC
- STACK
- PERCALL_COST
- SELCOST

Specifying the External Name

The following diagram details the external-name portion of the CREATE ROUTINE (or FUNCTION or PROCEDURE) statement for a UDR written in Java code.

►—EXTERNAL NAME—| Java External Name |—————►

Java External Name:



Element	Purpose	Restrictions
<i>class_name</i>	Class to which the UDR belongs	Must be an existing class.
<i>database</i>	Database where the jar exists If omitted, defaults to the current database.	Must be an existing database.
<i>jar_name</i>	Jar identifier as specified in the install_jar() statement	Must be an existing JAR name.
<i>java_datatype</i>	Name of a Java data type The second column of the following table shows data types and class names that you can use for this variable.	Must be a Java data type.
<i>method</i>	Name of the static method of the UDR	Must be an existing method.
<i>owner</i>	Owner of the jar If omitted, defaults to the current user.	Must be an existing user name.
<i>package_name</i>	Name of a package	Required if the UDR classes are in a package.

When used within a deployment descriptor, the THISJAR keyword automatically expands to the SQLJ-defined three-part JAR path.

The following table shows mapping between SQL data values and Java types. Use the values in the second column for the *java_datatype* variable.

SQL Data Type Java Type

CHAR(1)	char
CHAR(1)	java.lang.Character
CHAR()	Java.lang.String
CHARACTER()	java.lang.String
CHARACTER VARYING()	java.lang.String
VARCHAR	java.lang.String
LVARCHAR	java.lang.String
SMALLINT	short
SMALLINT	java.lang.Short
INTEGER	int
INTEGER	java.lang.Integer
INT8	long
INT8	java.lang.Long
SMALLFLOAT	float
SMALLFLOAT	java.lang.Float
REAL	float
REAL	java.lang.Float
FLOAT	double
FLOAT	java.lang.Double
DOUBLE PRECISION	double
DOUBLE PRECISION	java.lang.Double
DECIMAL	java.math.BigDecimal
MONEY	java.math.BigDecimal
NUMERIC	java.math.BigDecimal
BOOLEAN	boolean
BOOLEAN	java.lang.Boolean
DATE	java.sql.Date
DATETIME HOUR TO SECOND	java.sql.Time
DATETIME YEAR TO FRACTION	java.sql.Timestamp

INTERVAL	java.lang.String
BLOB	java.sql.Blob
CLOB	java.sql.Clob

Using a Deployment Descriptor

A *deployment descriptor* allows you to include in a JAR file the SQL statements for creating and dropping the UDRs. Both `sqlj.install_jar()` and `sqlj.remove_jar()` take parameters that, when set appropriately, cause the procedure to search for deployment descriptor files in the JAR file.

You can include the following SQL statements in a deployment descriptor:

- CREATE FUNCTION
- CREATE PROCEDURE
- GRANT
- DROP FUNCTION
- DROP procedure

When you execute `sqlj.install_jar()` or `sqlj.remove_jar()`, the database server automatically performs the actions described by any deployment-descriptor files that exist in the JAR file.

Warning: The transaction handling of the current database controls the SQL statements that the deployment descriptor executes. Use a BEGIN WORK statement to begin a transaction before you execute the `sqlj.install_jar()` or `sqlj.remove_jar()` procedure. In this way, a successful deployment can be committed, while a failed deployment can be rolled back.

For example, you might prepare a file, **deploy.txt**, that includes the following statements:

```
SQLActions[] = {
  "BEGIN INSTALL
    CREATE PROCEDURE showusers()
      WITH (class='jvp')
      EXTERNAL NAME 'thisjar:admin.showusers()'
      LANGUAGE JAVA;
    GRANT EXECUTE ON PROCEDURE showusers() to informix;
  END INSTALL",

  "BEGIN REMOVE
    DROP PROCEDURE showusers();
  END REMOVE"
```

For details on deployment-descriptor files, refer to the SQLJ: SQL Routines specification.

Using a Manifest File

The *manifest file* specifies the names of the deployment descriptor files that a JAR file contains. The **m** option of the **jar** command incorporates the manifest file into the default manifest of the JAR.

The following example shows the manifest file, **manifest.txt**, for a JAR with two deployment descriptors:

```
Name: deploy1.txt
SQLJDeploymentDescriptor: TRUE
```

```
Name: deploy2.txt
SQLJDeploymentDescriptor: TRUE
```

The following example shows the **jar** command that incorporates **manifest.txt** into a JAR file:

```
jar cvmf manifest.txt admin.jar deploy*.txt *.class
```

Compiling the Java Code

A UDR written in Java code is implemented by a static method in a Java class.

To make the Java source code into an executable format:

1. Compile the **java** files with the **javac** command to create class files.
2. Use the **jar** command to collect a set of class files into a JAR file.

For example:

```
# makefile for admin class
JAR_NAME = admin.jar
all:
    javac *.java
    jar cvmf manifest.txt $(JAR_NAME)
        deploy.txt *.class
    mv $(JAR_NAME) $(INFORMIXDIR)/jars
cleanup:
    rm -f *.class $(INFORMIXDIR)/jars/$(JAR_NAME)
```

JAR files contain Java classes that in turn contain static methods corresponding to SQL UDRs. JAR files can also contain auxiliary classes and methods that are used by the UDRs (for example, to perform SQL-to-Java type mapping).

Installing a JAR File

JAR files contain the code for the UDRs. For an SQL statement to be able to include a UDR written in Java code, you must install the jar file in the current database. Once a JAR file is installed, the routine manager of the database server can load the appropriate Java class when the UDR is invoked.

To manage jar files, use the EXECUTE PROCEDURE statement with the following SQLJ built-in procedures:

- **sqlj.install_jar(jar_url varchar(255), jar_id varchar(255), deploy_flag int)**
Before a Java static method can be mapped to a UDR, the class file that defines the method must be installed in the database. The **install_jar()** procedure installs a Java JAR file in the current database and assigns it a *jar identifier* (or *jar id*) for use in subsequent CREATE FUNCTION or CREATE PROCEDURE statements.

For example:

```
-- Installing a jar file
EXECUTE PROCEDURE sqlj.install_jar
('file:$INFORMIXDIR/jars/admin.jar',
 'admin_jar', 1);
```

- **sqlj.replace_jar(jar_url varchar(255), jar_id varchar(255))**
The **replace_jar()** procedure replaces a previously installed jar file with a new version.
- **sqlj.remove_jar(jar_id varchar(255), undeploy_flag int)**
The **remove_jar()** procedure removes a previously installed jar file from the current database.
- **sqlj.alter_java_path(jar_id varchar(255), path lvarchar)**
The **alter_java_path()** procedure specifies the *java-file search path* to use when the routine manager resolves related Java classes for the JAR file of a UDR.

For details about jar-naming conventions, refer to the SQLJ: SQL Routines specification.

All SQLJ built-in procedures reside in the **sqlj** schema.

Both **sqlj.install_jar()** and **sqlj.remove_jar()** take a parameter that, when set appropriately, causes the procedure to execute the deployment descriptor files in the JAR file.

For more information about how to install jar files, refer to the SQLJ: SQL Routines section of the documentation on the following Web site:

<http://www.sqlj.org/>

The SQLJ: SQL Routines specification has detailed tutorials on writing, registering, installing, and calling routines written in Java code.

Updating JAR Filenames

The script `update_jars.sql` is provided to update the three-part names of installed JAR files when you rename the database to which the JAR file belongs. You must execute the `update_jars.sql` script in the database after you rename it. You need to execute the `update_jars.sql` script only if you rename a database that has one or more installed JAR files.

Executing the User-Defined Routine

After you register a UDR as an external routine in the database, the UDR can be invoked in SQL statements such as:

- In the select list of a SELECT statement
- In the WHERE clause of a SELECT, UPDATE, or DELETE statement
- With the EXECUTE PROCEDURE or EXECUTE FUNCTION statement

The routine manager of the database server handles the execution of the UDR. For more information about the routine manager, see the *IBM Informix: User-Defined Routines and Data Types Developer's Guide*.

Debugging and Tracing

As with a UDR written in C, a UDR written in Java code might generate the SQL messages for UDR and DataBlade API errors when it executes. UDRs written in Java code adopt the JDBC error-reporting mechanism as well. The UDR throws an **SQLException** in case of an execution error such as a failed JDBC call. The routine manager detects such exceptions and translates it into a normal UDR error message.

In addition, the UDR can generate Java trace outputs and stack dumps at runtime. These additional Java messages are written to the *JVP log file*. The JVP log file is separate from the main database server log file, **online.log**. No JVP-specific messages appear in the database log. The JVP log file is intended to be the main destination for logging and tracing messages that are specific to the JVP and the UDR. This log is essential to support and debugging efforts. You should preserve it when possible.

Generating Log Messages

Log messages in the JVP log file can originate from any of the following sources:

- The JVP

JVP messages report such conditions as:

- JVP status (such as boot progress)
- Warnings about missing or limited resources
- Execution errors (such as being unable to locate a UDR)
- Internal errors (such as unexpected exceptions)

JVP log messages that report serious errors usually print a Java-method stack trace.

- The UDR

Log messages from the UDR are messages that make sense only in the JVP and Java domain or that can complement the messages from SQL or the database server with annotations and references that are specific to Java code or the JVP.

Use the following methods to write messages to the JVP log file from within a UDR:

- **UDRLog.log()**
- **UDRTraceable.tracePrint()**

By default, the JVP uses the following log file:

```
./jvp.log
```

where '.' is the current directory of the user who runs **oninit**.

You can change this default log file with the JVPLOGFILE parameter in the ONCONFIG configuration file. Set this configuration parameter to the name of the log file that you want the JVP to use. For example, the following line sets the log file to **/usr/jvp.log**:

```
JVPLOGFILE /usr/jvp.log
```

Important: Do *not* use the JVP log for error messages that need to be reported to the client application or to the main **online.log** file. Instead, the method should throw an **SQLException**.

Using the Administrative Tool

The IBM Informix JDBC Driver includes a built-in iterative UDR that is a limited administrative tool, **informix.jvpcontrol()**. The database server enables the **informix.jvpcontrol()** UDR when the JVPPROFILE configuration parameter specifies a starting port number by using the **JVP.monitor.port** entry.

You invoke **informix.jvpcontrol()** with the following syntax:

```
EXECUTE FUNCTION informix.jvpcontrol (command lvarchar);
```

The *command* can be one of the following forms, where *vpid* is the virtual processor ID:

- **threads** *vpid*
- **memory** *vpid*

You can use the **onstat -g glo** command to list the **vpid** numbers.

The threads **vpid** Option

The **threads vpid** form lists the threads running on the Java VP whose ID is **vpid**. For example, if *command* is threads 4, the UDR might return the following output:

```
(expression) Thread[informix.jvp.dbapplet.impl.JVPControl#0,
9,informix.jvp.dbapplet.impl.JVPControl#0],UDR=JVPControlUDR(java.lang.String), state = EXECUTE
(expression) Thread[JVP control monitor thread,10,main]
(expression) Thread[main,10,main]
(expression) Thread[SIGQUIT handler,0,system]
(expression) Thread[Finalizer thread,1,system]
5 row(s) retrieved.
```

The memory **vpid** Option

The **memory vpid** form lists memory use on the Java VP whose ID is **vpid**. For example, if *command* is memory 4, the UDR might return the following output:

```
(expression) Memory 16521840 bytes free, 16777208 bytes total
1 row(s) retrieved.
```

Debugging a Java User-Defined Routine

To debug a UDR written in Java code, you can connect the Java debugger, **jdb**, to the embedded JVM for debugging. The agent password that **jdb** requires is printed in the message log.

Traceable Events

The database server provides a fixed set of system trace events such as UDR sequence initialization, activation, and shutdown. You can also generate application-specific traces. For more information, refer to “The com.informix.udr.UDRTraceable” on page 4-7.

Finding Information about User-Defined Routines

The system catalog tables contain information about UDRs. The LANGUAGE clause of the CREATE FUNCTION or CREATE PROCEDURE statement tells the database server in which language the UDR is written. For UDRs in Java code, the LANGUAGE clause must be as follows:

```
LANGUAGE JAVA
```

The database server stores valid UDR languages in the **sysroutinelangs** table. The information includes an integer, the *language identifier*, in the **langid** column. The following lines show the entry in the **sysroutinelangs** system catalog table for the Java language:


```

langid      3
langname    java
langinitfunc  udr1m_java_init
langpath    $INFORMIXDIR/extend/krakatoa/1mjava.so
langclass   jvp

```

The Java language has the same default privilege as the C language. The following entry in the **syslangauth** system catalog table specifies the privileges for the Java language:

```

grantor      informix
grantee      DBA
langid       3
langauth     u

```

By default, both user **informix** and the owner of the database are allowed to create UDRs in Java code. If you attempt to execute the CREATE FUNCTION or CREATE PROCEDURE statement as some other user, the database server generates an error.

To allow other users to register UDRs in the database, user **informix** can grant the usage privilege on the Java language with the GRANT statement. The following GRANT statement allows any user who has Resource privileges on the database to register UDRs written in Java code:

```
GRANT USAGE ON LANGUAGE JAVA TO public
```

For more information on the syntax of the GRANT statement, see the *IBM Informix: Guide to SQL Syntax*.

Complying with SQLJ

The syntax of Java UDRs that the Informix database server supports usually follows the SQLJ specification. Where syntactic differences and missing features occur, the differences are mostly due to differences between Informix SQL and the SQL-3 standards. The following table summarizes the level of SQLJ compliance.

Feature (SQLJ Section #)	Function	Syntax	Definition and Rules	Comments
jar names (3.1)	Yes	Yes	Yes	
Java path (3.2)	Yes	Yes	Yes	
Install, replace, or remove jars (4.1-4.3)	Yes	Yes	Yes (required) No (optional)	No support of the optional replacement jar validation rules.
Alter java path (4.4)	Yes	Yes	Yes	

Feature (SQLJ Section #)	Function	Syntax	Definition and Rules	Comments
Create procedure, Create function (5.1)	Yes	Yes	Yes (required) No (optional)	No support of the optional create time jar validation and the Java main method.
For information about modifiers for Create Procedure and Create Function, refer to “Unsupported Modifiers” on page 4-20 and “Unsupported Optional Modifiers” on page 4-21.				
Drop procedure, Drop function (5.2)	Yes	Yes	Yes	
Grant or revoke jar (5.3-5.4, optional)	No	No	No	
SQLJ function call (5.5)	Yes	Yes	Yes	
SQLJ procedure call (5.6)	Yes	Yes	Yes	
System properties and default connections	No	No	No	
Deployment- descriptor files (optional)	Yes	No	No	
Status codes, exception handling (7.1-7.2)	Yes	Yes	Yes	

Unsupported Modifiers

Some modifiers for CREATE PROCEDURE and CREATE FUNCTION are not supported in this version of the database server. Informix UDRs do not support the following routine modifiers of the SQLJ specification.

Modifier	How to handle the modifier
Read SQL data	No Informix equivalent
Contains SQL	No Informix equivalent
Modifies SQL data	No Informix equivalent
No SQL	No Informix equivalent
Return null on null input	Informix default for external routines
Call on null input	Use the Informix modifier HANDLESNULLS
Deterministic	Use the Informix modifier NOT VARIANT

Nondeterministic	Use the Informix modifier VARIANT
Returns Java data type in Java method signature	No Informix equivalent
In parameter	Informix default; no need to specify the modifier

Unsupported Optional Modifiers

Informix UDRs do not support the following optional routine modifiers of the SQLJ specification:

- Dynamic result sets
- Inout parameter
- Output parameters in callable statements

Chapter 5. The IBM Informix JDBC Driver

Public JDBC Interfaces	5-1
The <code>com.informix.jdbc.IfxConnection</code>	5-2
The <code>com.informix.jdbc.IfxProtocol</code>	5-2
The <code>informix-direct</code> Subprotocol	5-2
JDBC 1.0 API	5-3
JDBC 2.0.	5-4
Support for Opaque Data Types	5-5
<code>java.sql.SQLUDTInput</code>	5-5
<code>java.sql.SQLUDTOutput</code>	5-5
Interfaces Updated for Java 2.0	5-5
An Example That Shows Query Results	5-5

In This Chapter

All UDRs written in Java code can access the database server data through the JDBC application programming interface (API). This chapter briefly describes the Informix implementation of the JDBC API and the server-side IBM Informix JDBC Driver.

Generally, the IBM Informix server-side JDBC driver derives from the client-side driver so that the two drivers are essentially the same. Java UDRs require some differences, however, to use the IBM Informix JDBC Driver from the server side. This chapter describes the public JDBC interfaces and JDBC subprotocols that the IBM Informix JDBC Driver provides specifically for server-side JDBC applications, as well as restrictions that apply to server-side JDBC applications. For principal documentation of the IBM Informix JDBC Driver, refer to the *IBM Informix: JDBC Driver Programmer's Guide*.

Public JDBC Interfaces

IBM Informix JDBC Driver defines the following public interfaces:

- `com.informix.jdbc.IfxConnection`
- `com.informix.jdbc.IfxProtocol`

The client and server JDBC drivers each have their own implementation of the preceding interfaces. The client driver provides access to databases from Java applications. The server driver provides database access from within the server through UDRs written in Java code.

The `com.informix.jdbc.IfxConnection`

The `IfxConnection` interface is a subinterface of `java.sql.Connection` with Informix-specific methods added. The `com.informix.jdbc.IfxDirectConnection` class implements the `com.informix.jdbc.IfxConnection` interface. This interface provides a connection to the current database server from within a UDR. The connection corresponds to a server-query context and is passed to the UDR by the SQLJ language manager. The transaction context of this connection is that of the query issuing the UDR call, and the call to create a UDR connection does not specify any database or user information.

The `com.informix.jdbc.IfxProtocol`

The `IfxProtocol` interface represents the protocol and data exchange between the client application and an Informix database server. It sends and processes the messages and data flow between the client and database server. The `com.informix.jdbc.IfxDirectProtocol` class implements the `IfxProtocol` interface. It uses the DataBlade API (DAPI) to access database resources.

The `informix-direct` Subprotocol

The JDBC `DriverManager` class provides services to connect to JDBC drivers. It assists in loading and initializing a requested JDBC driver. A UDR written in Java code uses the `registerDriver()` method of `DriverManager` to register itself and to redirect user messages to the `DriverManager` logging facility.

A UDR written in Java code or a Java client application that wants to connect to the database calls the `DriverManager.getConnection()` method to obtain a connection handle. This method takes a URL string as an argument. The JDBC management layer attempts to locate a driver that can connect to the database that the URL represents. To perform this task, the JDBC management layer asks each driver in turn if it can connect to the specified URL. Each driver examines the URL and determines if it supports the specified JDBC subprotocol. The Informix implementation of UDRs written in Java code supports the `informix-direct` subprotocol in the database server.

For the `informix-direct` subprotocol, the JDBC driver loads and uses the following classes:

- The *connection class*, which you can specify with the `ConnectionClass` property. The connection class must implement `IfxConnection`.
- The *protocol class*, which you can specify with the `ProtocolClass` property. This protocol class must implement `IfxProtocol`.

These specifiers are optional in the URL string. If you do not specify `ConnectionClass` or `ProtocolClass`, the IBM Informix JDBC Driver can determine them from the subprotocol.

The following call opens a UDR connection with the class **IfxDirectConnection**. It uses the **IfxDirectProtocol** as the protocol for processing queries on the current database.

```
DriverManager.getConnection("jdbc:informix-direct:"+  
"//ConnectionClass=com.informix.jdbc.IfxDirectConnection;" +  
"//ProtocolClass=com.informix.jdbc.IfxDirectProtocol");
```

The UDR connection can only be opened by the thread that executes the UDR static method. In this way, the database server can ensure that the proper transaction context is used for the UDR.

JDBC 1.0 API

The JDBC 1.0 API consists of the following Java classes and interfaces that you can use to open connections to particular databases, execute SQL statements, and process the results.

Classes	Interfaces
java.sql.DataTruncation	java.sql.CallableStatement
java.sql.Date	java.sql.Connection
java.sql.DriverManager	java.sql.DatabaseMetaData
java.sql.DriverPropertyInfo	java.sql.Driver
java.sql.SQLException	java.sql.PreparedStatement
java.sql.SQLWarning	java.sql.ResultSet
java.sql.Time	java.sql.ResultSetMetaData
java.sql.Timestamp	java.sql.Statement
java.sql.Types	None

The following JDBC 1.0 classes and interfaces are the most important for the development of UDRs in Java code:

- **java.sql.DriverManager** handles loading of drivers and provides support for creating new database connections.
- **java.sql.Connection** represents a connection to a particular database.
- **java.sql.Statement** acts as a container for executing an SQL statement on a given connection.
- **java.sql.ResultSet** controls access to the row results of a given statement.
- **java.sql.PreparedStatement** handles execution of a pre-compiled SQL statement.
- **java.sql.CallableStatement** handles execution of a call to a database SPL routine.

For more documentation, refer to the JavaSoft Web site at:
<http://java.sun.com>

JDBC 2.0

JDBC 2.0 is a major leap from JDBC 1.0 in that it supports extensible data types and large objects. The following extensions to JDBC 1.0 are provided to support user-defined data types (UDTs) with JDK 1.1.x:

- **java.sql.Blob**
- **java.sql.Clob**
- **java.sql.SQLData**
- **java.sql.SQLInput**

The following read/write methods are not supported for opaque types:

- **readString()**
Use the Informix extension **readString(len)**.
- **readInterval()**
- **readBytes()**
Use the Informix extension **readBytes(len)**.
- **readCharacterStream()**
- **readAsciiStream()**
- **readBinaryStream()**
- **readObject()**
- **readRef()**
- **readArray()**

- **java.sql.SQLOutput**

The following read/write methods are not supported for opaque types:

- **writeString()**
Use the Informix extension **writeString(len)**.
- **writeInterval()**
- **writeBytes()**
Use the Informix extension **writeBytes(len)**.
- **writeCharacterStream()**
- **writeAsciiStream()**
- **writeBinaryStream()**
- **writeObject()**
- **writeRef()**
- **writeArray()**

Support for Opaque Data Types

Certain JDBC 2.0 interfaces need to be extended to support opaque data types. Some of the methods need an additional length argument to read or write an opaque data type because the JDBC driver cannot look inside an opaque data type to determine the field lengths.

The Informix implementation of UDRs written in Java code provides the following extensions of the JDBC user-defined-type (UDT) support:

- **java.sql.SQLUDTInput**
- **java.sql.SQLUDTOutput**

For more information on using an opaque data type in a Java UDR, refer to Chapter 6, “Using Opaque User-Defined Types,” on page 6-1.

java.sql.SQLUDTInput

This class extends **java.sql.SQLInput** with the following methods:

```
public String readString(int maxlen) throws SQLException;
public byte[] readBytes(int maxlen) throws SQLException;
```

java.sql.SQLUDTOutput

This class extends **java.sql.SQLOutput** with the following methods:

```
public void writeString(String str, int maxlen) throws SQLException;
public void writeBytes(byte[] b, int maxlen) throws SQLException;
```

Interfaces Updated for Java 2.0

The Informix implementation of UDRs written in Java code also defines the following public interfaces:

- **com.informix.PreparedStatement2**
This class includes the JDBC 2.0 methods **setBlob()** and **setClob()**.
- **com.informix.ResultSet2**
This class includes the JDBC 2.0 methods **getBlob()** and **getClob()**.
- **com.informix.Types2**
This class includes the type codes for the smart-large-object data types, BLOB and CLOB.

An Example That Shows Query Results

The following example implements a procedure called **showusers()**, which runs a query, retrieves all rows from the returned result, and prints the rows in the JVP log file:

```
import com.informix.udr.*;
import java.sql.*;

public class admin
```

```

{
    public static void showusers() throws SQLException
    {
        UDREnv env = UDRManager.getUDREnv();
        UDRLog log = env.getLog();
        String name = env.getName();

        Connection conn = DriverManager.getConnection
            ("jdbc:informix-direct:");
        Statement stmt = conn.createStatement();
        ResultSet rs = stmt.executeQuery
            ("SELECT * FROM Users");
        log.log("User information:");

        while ( rs.next() )
        {
            String UID = rs.getString(1);
            String Password = rs.getString(2);
            String Last = rs.getString(3);
            String First = rs.getString(4);

            // Write out the UDR name followed by the
            // columns values
            String line = name + " : " +
                UID + " " + Password + " " + Last + " " + First;
            log.log(line);
        }
        stmt.close();
        conn.close();
    }
}

```

After you create and install the JAR file that contains this Java method, the next task is to register the **showusers()** method as a UDR by giving it an SQL procedure signature. For the CREATE PROCEDURE statement that registers **showusers()**, see “Specifying the JVP” on page 4-9.

The syntax for invoking a UDR written in Java code is no different from a standard UDR call, as follows:

```
EXECUTE PROCEDURE showusers()
```

Chapter 6. Using Opaque User-Defined Types

Using the <code>SQLData</code> Interface	6-1
Default Input/Output Routines	6-2
SQL Definitions for Default I/O User-Defined Routines	6-3
Informix Extensions to <code>SQLInput</code> and <code>SQLOutput</code>	6-3
<code>IfmxUDTSQLInput</code>	6-3
<code>IfmxUDTSQLOutput</code>	6-4
The Circle Class Example	6-5
Overriding the Default I/O Methods	6-6
I/O Function Sets and Related Types	6-6
<code>IfxDatapointer</code>	6-7
Stream Implementations	6-8
An Example That Overrides the Default I/O Methods	6-9
Usage Example	6-10
Limitations to Streams.	6-15

In This Chapter

This chapter describes how to use opaque user-defined data types (UDTs). It describes the default `SQLData` interface, as well as how to override the default. It provides the following information:

- The `SQLData` Interface
- SQL statements to create default I/O routines
- Informix extensions to `SQLInput` and `SQLOutput` interfaces
- How to override the default I/O methods
- Required I/O function sets and related data types
- Limitations to Streams

Using the `SQLData` Interface

To implement a complete UDT in Java code, you must supply a set of data-formatting methods that convert to and from the various representations of the data type. These methods perform input and output operations for the data type such as converting text input to the internal structure that the database server uses.

All the database server I/O functions manipulate data formats that can be represented as Java streams. The streams encapsulate the data and implement methods needed to parse the source format or write the destination format.

To implement an opaque UDT and use the default data-translation I/O methods:

1. Supply the JDBC `SQLData` interface: `readSQL()`, `writeSQL()`, and `getSQLTypeName()` methods.
2. Create the SQL routine and cast definitions for the I/O functions by calling `sqlj.registerJUDTfuncs(vchar(255))`, where the `vchar` argument is the SQL name of the type you are registering.

For example, after creating the UDT `Record3` with the following statements:

```
create opaque type Record3 (internallength = variable,  
    alignment = 8, maxlen = 2048, cannothash );  
grant usage on type Record3 to public;  
execute procedure setUDTExtName("Record3",  
    "informix.testclasses.jlm.udt.Record3");
```

you could create the default casts and I/O functions with the following statement:

```
execute procedure registerJUDTfuncs("Record3");
```

The `readSQL()` method converts a database type to a Java object and the `writeSQL()` method converts a Java object to the database type. The system supplies the appropriate stream type at runtime.

To back out default I/O methods for an opaque UDT:

You can back out default I/O functions and casts by calling `sqlj.unregisterJUDTfuncs(vchar(255))`, where the `vchar` argument is the SQL name of the type, as the following example shows:

```
execute procedure unregisterJUDTfuncs("Record3");
```

Default Input/Output Routines

Because this interface uses Java, all the SQL I/O support functions are predefined when you register the UDT. You only need to supply the required `SQLData` implementation.

Informix supplies extensions to the Stream arguments of `SQLData` methods to suit various uses. With these extensions, you can build I/O functions for a new Java UDT. All that you must do to implement any of the required function sets is select the Stream type.

Informix also supplies default Input and Output processing methods in Java code that are used to implement all UDT I/O operations. The database server contains these default I/O methods and executes them just like any other Java UDR. These methods use information in the SQL UDR definition to select the correct Streams and instantiate the right user-defined type objects at execution time.

“The Circle Class Example” on page 6-5 illustrates the use of the **SQLData** interface.

SQL Definitions for Default I/O User-Defined Routines

After you register the Java UDT with the database server using the SQL procedure **setUDTExtName()**, you can create SQL functions and casts for it, using either the default I/O wrapper methods or explicit methods in your Java UDT class. For the default I/O wrapper methods, the **registerJUDTfuncs** function creates the SQL functions shown in the following example, where **SQLType** is the SQL UDT name, **JavaType** is the JUDT name, and **SQLBuffer** is the SQL *transport type* being converted, that is, SENDRECV:

```
-- Receive function

CREATE IMPLICIT CAST (SENDRECV as SQLUDT with
    IfxJavaSENDRECVInJavaUDT);
CREATE FUNCTION IfxJavaSENDRECVInJavaUDT (in SENDRECV)
    RETURNS SQLUDT
EXTERNAL NAME
'com.informix.jdbc.IfxDatapointer.IfxDatapointer(java.lang.Object)'
LANGUAGE java;
GRANT EXECUTE ON FUNCTION IfxJavaSENDRECVInJavaUDT TO PUBLIC;

-- Send function

CREATE EXPLICIT CAST (SQLUDT as SENDRECV with
    IfxJavaSENDRECVOutJavaUDT);
CREATE FUNCTION IfxJavaSENDRECVOutJavaUDT(out SQLUDT) RETURNS
    SENDRECV
EXTERNAL NAME 'com.informix.jdbc.IfxDatapointer.IfxDatapointer(java.sql.SQLData)'
LANGUAGE java NOT VARIANT;
GRANT EXECUTE ON IfxJavaSENDRECVOutJavaUDT TO PUBLIC;
```

The default Input method cannot be declared *not variant* because it might need to perform SQL queries to instantiate the correct Java UDT class.

Informix Extensions to SQLInput and SQLOutput

Some of the standard **SQLInput** and **SQLOutput** Stream methods need an additional length argument to read or write an opaque data type because the JDBC driver cannot determine the field lengths for an opaque type. Informix database server provides the **IfmxUDTSQLInput** and **IfmxUDTSQLOutput** extensions, which inherit from the standard JDBC 2.0 **SQLInput** and **SQLOutput** interfaces.

IfmxUDTSQLInput

The **IfmxUDTSQLInput** interface extends **SQLInput**, which contains the following public methods:

```
String readString()
boolean readBoolean()
byte readByte()
short readShort()
```

```

int readInt()
long readLong()
float readFloat()
double readDouble()
java.math.BigDecimal readBigDecimal()
byte[] readBytes()
java.sql.Date readDate()
java.sql.Time readTime()
java.sql.Timestamp readTimestamp()
java.io.Reader readCharacterStream()
java.io.InputStream readAsciiStream()
java.io.InputStream readBinaryStream()
Object readObject()
Ref readRef()
Blob readBlob()
Clob readClob()
Array readArray()
boolean wasNull()

```

The **IfmxUDTSQLInput** interface adds the following Informix methods:

```

String readString(int maxlen)
byte[] readBytes(int maxlen)
Interval readInterval()
int available();
int length();
IfxUDTInfo getUDTInfo(int xid)
IfxUDTInfo getUDTInfo(String name, String owner)

```

All the **readXXX()** methods throw an **SQLException** when they detect parsing errors. Use the **readXXX()** methods to convert the buffer of the given Input stream into a Java object. When the Input stream is empty, each read method throws an **SQLException** with **e.getErrorCode** equal to **-79772** or **IfxErrMsg.S_BADSQLDATA**. However, you can use the **length()** and **available()** methods to determine when the Input stream is exhausted while converting variable length UDTs to Java objects.

IfmxUDTSQLOutput

The **IfmxUDTSQLOutput** interface extends **SQLOutput**, which contains the following public methods:

```

void writeString(String x)
void writeBoolean(boolean x)
void writeByte(byte x)
void writeShort(short x)
void writeInt(int x)
void writeLong(long x)
void writeFloat(float x)
void writeDouble(double x)
void writeBigDecimal(java.math.BigDecimal x)
void writeBytes(byte[] x)
void writeDate(java.sql.Date x)
void writeTime(java.sql.Time x)
void writeTimestamp(java.sql.Timestamp x)

```

```

void writeCharacterStream(java.io.Reader x)
void writeAsciiStream(java.io.InputStream x)
void writeBinaryStream(java.io.InputStream x)
void writeObject(SQLData x)
void writeRef(Ref x)
void writeBlob(Blob x)
void writeClob(Clob x)
void writeStruct(Struct x)
void writeArray(Array x)

```

The **IfmxUDTSQLOutput** interface adds the following Informix methods:

```

void writeString(String x, int length)
void writeBytes(byte[] b, int length)
void writeInterval(Interval intrvl)
int available()
int length()
IfxUDTInfo getUDTInfo(int xid)
IfxUDTInfo getUDTInfo(String name, String owner)

```

All the **writeXXX()** methods throw an exception when they encounter conversion errors. Use the Stream **write()** methods to convert a Java object into the given Output buffer. The **length()** method returns the number of bytes that remain in the buffer. The JDBC 2.0 class files describe the **SQLOutput** definition.

The Circle Class Example

The **circle** class example implements a fixed-length opaque data type. The **circle** data type includes X and Y coordinates (**xCoord** and **yCoord**), which represent the center of the circle and a radius value (**radius**). The **readSQL** method reads the input stream **SQLInput** to obtain the **xCoord**, **yCoord**, and **radius** values and saves the data type name from String **typename**. The **writeSQL** method writes the **xCoord**, **yCoord**, and **radius** values to the stream **SQLOutput**.

```

package informix.testclasses.jlm;

import java.sql.*;

public class circle implements SQLData
{
    public int xCoord;
    public int yCoord;
    public int radius;
    private String type;

    public String getSQLTypeName()
    {
        return type;
    }

    public void readSQL (SQLInput stream, String typeName)
        throws SQLException

```

```

    {
        xCoord = stream.readInt();
        yCoord = stream.readInt();
        radius = stream.readInt();

        type = typeName;
    }

    public void writeSQL (SQLOutput stream)
        throws SQLException
    {
        stream.writeInt(xCoord);
        stream.writeInt(yCoord);
        stream.writeInt(radius);
    }
}

```

The **SQLData** methods use I/O streams to translate between C and Java representations. The following C-language structure shows the C definition for the circle:

```

typedef struct
{
    int x;
    int y;
    int radius;
} circle;

```

Overriding the Default I/O Methods

If the default methods are not sufficient because, for example, you want to include parentheses and other delimiting characters in the text representation, you can explicitly override the defaults with definitions of your own, after you register the Java UDT.

I/O Function Sets and Related Types

Table 6-1 specifies the I/O functions that you must implement for the nondefault case, and their related data types.

Table 6-1. Nondefault I/O Functions and Types Table

Function Set	Data Format	Buffer Type		Java Stream Implementation
		SQL	Java	
Server UDR	UDT	Internal Representation	IfxDatapointer	IfmXSQLInStream IfmXSQLOutStream
Input Output	Text	LVARCHAR	String (String Buffer)	IfmXTextInStream IfmXTextOutStream
Send Receive	Client Binary	SENDRECV	IfxDatapointer	IfmXSRInStream IfmXSROutStream

Table 6-1. Nondefault I/O Functions and Types Table (continued)

Function Set	Data Format	Buffer Type		Java Stream Implementation
		SQL	Java	
Import Export	Text	IMPEXP	IfxDatapointer	IfmxIEInStream IfmxIEOutStream
Binary Import Export	Client Binary	IMPEXBIN	IfxDatapointer	IfmxIEBinStream IfmxIEBOutStream

The columns in the preceding table represent the following:

- Function set
Names the type of function in conformance with UDT specifications
- Data format
A conceptual description of the format of the data in the SQL buffer that is being converted
- Buffer type
Names the actual data types being read or written
 - SQLBuffer is the SQL (or database-server) type for this data.
 - JavaBuffer is the Java type to which the SQLBuffer is transformed prior to being passed to (or returned from) the I/O method.
It is an intermediate type that is contained in and manipulated by a Java Stream. It is also the argument type for input methods and the return type for output methods.
- Java Stream implementation
Names the actual stream type that is passed to the **SQLData** interface when the default I/O functions are used. Each of the streams implements **IfmxUDTSQLInput** or **IfmxUDTSQLOutput**.

IfxDatapointer

The **IfxDatapointer** class encapsulates the Informix C-language representation of a type and its corresponding data buffer. This is usually a database server buffer structure, with a few attributes extracted for easy access in Java code. This class is used to *transport* the nontextual SQL data types to and from the I/O methods and is generally managed by an **IfmxUDTSQLInput** or **IfmxUDTSQLOutput** stream.

Methods in both streams might throw an **SQLException** with the **e.getErrorCode** equal to **-79700** or **IfxErrMsg.S_MTHNSUPP**, if they are not implemented. These methods are generally not needed on the database server side but are useful in the client JDBC code.

For more documentation of these streams, refer to the *IBM Informix: JDBC Driver Programmer's Guide*. For an example of using these streams, see "Usage Example" on page 6-10.

Stream Implementations

The following sections briefly describe the Java classes that implement the **IfmxUDTSQLInput** and **IfmxUDTSQLOutput** interfaces.

IfmxSQLInStream and IfmxSQLOutStream: These streams convert to and from the internal data representation that the database server uses.

IfmxTextInStream and IfmxTextOutStream: These streams convert to and from a textual data representation for Input and Output functions. IBM Informix Dynamic Server with J/Foundation does not support cross-locale Input and Output routines; all strings are assumed to be in U.S. English.

These streams delimit each component of the composite type with a white space between record elements. The SQL type is an LVARCHAR that contains client text. The JavaBuffer type for Input is String, which contains the client text. The JavaBuffer type for Output is a StringBuffer. The **read()** and **write()** methods must convert between the client text representation and the relevant Java object.

IfmxSRInStream and IfmxSROutStream: These streams convert to and from the binary data representation of the client for send and receive functions. The SQL type is SENDRECV, which is an internal representation that contains binary data in the client format. The JavaBuffer type is **IfxDataPointer**. The **read()** and **write()** methods convert between the client representation and the relevant Java object.

IfmxIEInStream and IfmxIEOutStream: This stream converts to and from a canonical text representation for import and export functions. The **SQLBuffer** is an IMPEXP type that is an internal representation that contains canonical textual data. The JavaBuffer type is **IfxDataPointer**. The **read()** and **write()** methods convert between the text representation and the relevant Java objects. These streams inherit from the **IfmxTextInStream** and **IfmxTextOutStream** classes.

IfmxIEBInStream: This stream converts to and from a canonical binary representation for binary import and export functions. The **SQLBuffer** is an IMPEXPBIN type that is an internal representation that contains canonical binary data. The JavaBuffer type is **IfxDataPointer**. The **read()** and **write()** methods must convert between the binary representation and the relevant Java objects. These streams inherit from the **IfmxSRInStream** and **IfmxSRIOutStream** classes.

Class Layout (for Input): Figure 6-1 describes the class layout for input. The class layout for output is similar; simply replace In with Out in the names.

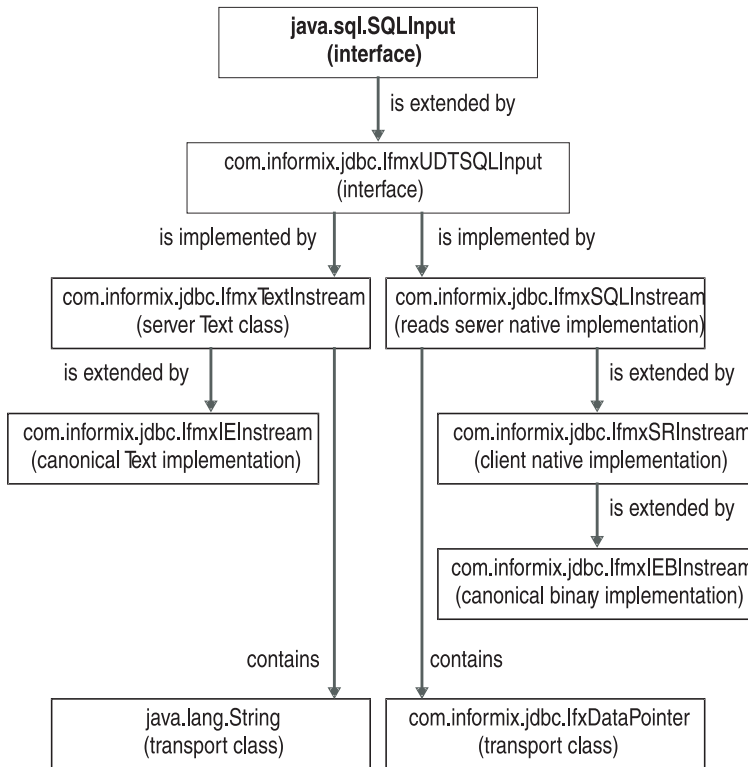


Figure 6-1. .Input Class Layout

An Example That Overrides the Default I/O Methods

The following example illustrates a Java UDT class with nondefault definitions. **JavaType** is the new Java UDT, and **JavaBuffer** is the buffer type for the SQL data being converted, as “I/O Function Sets and Related Types” on page 6-6 shows. For a complete set of required and optional code, see “Usage Example” on page 6-10.

```

public class JavaType implements SQLData
{
    // Java data Object declarations for this Class....
    // non-default Data Input function
    public static JavaType JavaTypeInput( JavaBuffer in )
    {
        JavaType x = new JavaType(); // make a new object
        // convert JavaBuffer fields to Java data objects in
        // this Class
        return( x );// return the new object
    }
}
  
```

```

// non-default Data Output function
public static JavaBuffer JavaTypeOutput( JavaType out )
{
    JavaBuffer x = new JavaBuffer();
    // Do whatever it takes to translate object to output
    // buffer format
return x;    // return the initialized buffer
}
// required SQLData implementation
private String type;
public String getSQLTypeName()
{
    return type;
}
public void readSQL ( SQLInput instream, String typeName )
    throws SQLException
{
    type = typeName;
    // cast up to Informix specific stream type
    IfmxUDTSQLInput in = (IfmxUDTSQLInput) instream;
    // read stream fields into Java data objects in this Class
    return;
}

public void writeSQL( SQLOutput outstream ) throws SQLException
{
    // cast up to Informix specific stream type
    IfmxUDTSQLOutput out = (IfmxUDTSQLOutput) outstream;
    // write object to output stream
    return;
}
}

```

For an example of the SQL definitions required to use the explicit methods in the preceding code, see “SQL Definitions for a Variable-Length UDT Example” on page 6-14.

Usage Example

All Java UDT classes must implement the **readSQL()** and **writeSQL()** methods for the **SQLData** interface. The **readSQL()** method initializes a Java object using data from the database server in a C-language format. The **writeSQL()** method converts a Java object back to the representation of the database server. The **readSQL()** and **writeSQL()** methods receive a Stream argument that encapsulates the conversion methods for each built-in type that the database server uses, for example, **int**, **float**, **decimal**.

In the case of a fixed-length UDT, the **readSQL()** and **writeSQL()** methods know the order and number of fields they are to process. In the case of a variable-length UDT, the programmer must rely on the **stream.available()** method and/or the **SQLException** to find the end of the data as this example shows.

Variable-Length UDT Including Nondefault Input and Output Methods:

```
/* Variable Length UDT example type: Record3
** Example of required and explicit method implementations.
**
** The C language structure equivalent of this JUDT is:
**
** typedef struct
** {
**     mi_double_precision d;
**     mi_chara[4];
**     mi_integerb;
**     mi_realc;
**     mi_datee;
**     mi_smallintf;
**     mi_booleang[MAXBOOLS];
** } NewFixUDT;
**
** Where the last boolean array can contain up to MAX values
** but only valid values will be written to disk.
*/
// Put this in our test package,
// could be anywhere but needs to match SQL definitons for UDRs.
package informix.testclasses.jlm.udt;
// get the usual suspect classes
import java.sql.*;
// get informix specific interfaces, etal.
import com.informix.jdbc.*;
// These are only needed for the non-default Input/Output
// functions, remove if you use defaults.
import informix.jsp.dbapplet.impl.IfmxTextInStream;
import informix.jsp.dbapplet.impl.IfmxTextOutputStream;
/***** Now here's our UDT *****/
public class Record3 implements SQLData
{
    // to turn debug print lines on and off
    private static boolean classDebug = true;

    // define storage for Java members of UDT
    private double d_double;
    private String a_char;
    private int b_int;
    private float c_float;
    private java.sql.Date e_date;
    private short f_smint;
    // could use a Vector for booleans, but would then need Boolean
    // objects ...so I've left it as an exercise for the reader...
    private static final int MAXBOOLS = 20;
    private boolean g_boolvals[] = new boolean[MAXBOOLS];
    private int numbools = 0;
    // dummy constructor just so we can log instantiation
    public Record3()
    {
        super();
        if( classDebug )
            System.out.println( "Record3() " + super.toString() + " created" );
    }
}
```

```

    }
    // dummy finalizer just so we can log our own destruction
    protected void finalize()
    {
        super.finalize();
        if ( classDebug )
            System.out.println( "Record3() " + super.toString() + " deleted" );
    }
    /***** REQUIRED SQLData implementation: *****/
    // needed for SQLData interface
    private String type;
    public String getSQLTypeName()
    {
        return type;
    }
    // Called to convert an SQL buffer TYPE to JAVA class.
    // note: we need to use SQLInput as the argument type or this
    // method signature won't resolve correctly.
    public void readSQL (SQLInput in, String typeName) throws
    SQLException
    {
        if( classDebug )
            System.out.println( "Record3.readSQL() entered" );
        // save the type name
        type = typeName;
        // cast the _real_ type of Stream for IFMX extensions.
        IfmxUDTSQLInput stream = (IfmxUDTSQLInput) in;
        // trap exceptions; don't really know how many bytes
        // are in the input.
        try
        {
            d_double = stream.readDouble();
            a_char = stream.readString(4);
            b_int = stream.readInt();
            c_float = stream.readFloat();
            e_date = stream.readDate();
            f_smint = stream.readShort();
            // Read booleans until we get an exception:
            // converting a non-existent boolean will throw cookies.
            // but we can use available() to make sure there is more
            // to read...
            for( int count = 0; (stream.available() > 0) && (count
                < MAXBOOLS); ++count )
            {
                g_boolvals[count] = stream.readBoolean();
                ++numbools;
            }
        }
        catch (SQLException e)
        {
            // if we got something besides end of input rethrow,
            // otherwise just assume we're done.
            if( e.getErrorCode() != IfxErrMsg.S_BADSQLDATA )
            {
                if( classDebug )

```

```

        System.out.println("Record3.readSQL() exception = " +
                           e.toString());
    }
    throw e;
}
}

// Called to convert JAVA class to SQL buffer TYPE.
// note: we need to use SQLOutput as the argument type or this
// method signature won't resolve correctly.

public void writeSQL( SQLOutput out ) throws SQLException
{
    if( classDebug )
        System.out.println( "Record3.writeSQL() entered" );
    // cast up to _real_ type of Stream to use IFMX extensions.
    IfmxUDTSQLOutput stream = (IfmxUDTSQLOutput) out;
    stream.writeDouble(d_double);
    stream.writeString(a_char, 4);
    stream.writeInt(b_int);
    stream.writeFloat(c_float);
    stream.writeDate(e_date);
    stream.writeShort(f_smint);
    for( int i = 0; i < numbools; i++ )
        stream.writeBoolean(g_boolvals[i]);
}
/***** END SQLData implementation *****/
/**** NON-DEFAULT implementation of Input and Output functions ****/
/* Remove all this if you only use the Defaults */

```

The following example illustrates the implementation of user-defined input and output functions that override the default I/O methods. If you use the default methods, you do not need to implement overriding methods like those that follow:

```

// Called as Input function to convert SQL lvarchar to JAVA class
public static Record3 fromString( String str )
{
    if( classDebug )
        System.out.println( "Record3.fromString(String) entered" );
    // Make a stream of the right kind.
    IfmxTextInStream stream = new IfmxTextInStream(str);
    // Make a new Java object of the right type.
    Record3 record = new Record3();
    // Just call readSQL ourselves.
    // For a real implementation you would probably copy all the
    // readXXX()'s and intersperse delimiting chars as needed...
    try
    {
        readSQL( stream, "Record3" );
    }
    catch (Exception e)
    {
        System.err.println(e.getMessage());
    }
}

```

```

        return record;
    }

    // Called as Output function; convert JAVA class to SQL lvarchar.
    // note: could use toString() directly,
    // except that the UDR method must be "static", and
    // it needs to take a Record3 as an argument....

    public static String makeString(Record3 x)
    {
        if( classDebug )
            System.out.println( "Record3.makeString() entered" );
        return x.toString();
    }

    // Might as well implement the standard toString() as long as
    // we're doing non-defaults. If a different method name is
    // used here, Object.toString() will be called when the class
    // gets printed out in debug lines....

    public String toString()
    {
        // Need to use a StringBuffer because we can't pass a
        // reference to a String to be initialized.
        // We could optimize by guessing at size of buffer, too.
        // StringBuffer str = new StringBuffer();
        // IfmxTextOutputStream stream = new IfmxTextOutputStream(str);
        // Just call writeSQL.
        // For a real implementation you would probably copy all the
        // writeXXX()'s and intersperse delimiting chars as needed...
        try
        {
            writeSQL( stream );
        }
        catch (Exception e)
        {
            System.err.println(e.getMessage());
            // not sure if we need to clear out result string?
            str.setLength(0);
        }
        return str.toString();
    }
}

```

SQL Definitions for a Variable-Length UDT Example: The SQL definitions for this example are:

```

-- VarLen UDT and support functions -----
create opaque type Record3 (internallength = variable,
    alignment = 8, maxlen = 2048, cannotohash );
grant usage on type Record3 to public;
-- register JUDT implementation...
-- note package name needs to match class file package
execute procedure setUDTExtName("Record3",
    "informix.testclasses.jlm.udt.Record3");
-- Definitions for NON_DEFAULT Input/Output functions.

```



```

-- this overrides the defaults setup above
-- LVARCHAR INPUT
drop cast (Record3 as lvarchar);
create implicit cast (Record3 as lvarchar with record3_output);
create function record3_input (l lvarchar) returns Record3
    external name
'informix.testclasses.jlm.udt.Record3.fromString(java.lang.String)'
    language java not variant;
grant execute on function record3_input to public;
-- CHAR INPUT
drop cast (Record3 as char(100));
create implicit cast (Record3 as char(100) with record3_rout);
create function record3_rin (c char(100)) returns Record3
    external name
'informix.testclasses.jlm.udt.Record3.fromString(java.lang.String)'
    language java not variant;
grant execute on function record3_rin to public;

-- LVARCHAR OUTPUT
drop cast (lvarchar as Record3);
create explicit cast (lvarchar as Record3 with record3_input);
create function record3_output (c Record3) returns lvarchar
    external name
'informix.testclasses.jlm.udt.Record3.makeString(informix.testclasses.jlm.udt.Record3)'
    language java not variant;
grant execute on function record3_output to public;
-- CHAR OUTPUT
drop cast (char(100) as Record3);
create explicit cast (char(100) as Record3 with record3_rin);
create function record3_rout (c Record3) returns varchar(100) external name
'informix.testclasses.jlm.udt.Record3.makeString(informix.testclasses.jlm.udt.Record3)'
    language java not variant;
grant execute on function record3_rout to public;

-- END definitions for NON_DEFAULT Input/Output functions.
-- end VarLen UDT and support functions -----
-- Example Usage ---
create table rec3tab (record_col Record3);
insert into rec3tab values ('665.999 JAVA 398 197.236 1952-04-10 47 f t t');
insert into rec3tab values ('667.000 Jive 983 791.632 2002-04-11 42 f f f f');
select * from rec3tab;

```

Limitations to Streams

The following limitations apply to the I/O streams in IBM Informix Dynamic Server with J/Foundation:

- BLOBs and CLOBs are not supported.
- Text Input and Output across locales is not supported.
- Text Input and Output for intervals is not supported.
- Time stamps are only supported in their full format. Qualifiers are not supported.

- Byte arrays, `byte[]`, and Object/Stream I/O are not supported for either text or binary operations.

Appendix. Accessibility

The syntax diagrams in the HTML version of this manual are available in dotted decimal syntax format, which is an accessible format that is available only if you are using a screen reader.

Dotted Decimal Syntax Diagrams

In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), the elements can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read punctuation. All syntax elements that have the same dotted decimal number (for example, all syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, the word or symbol is preceded by the backslash (\) character. The * symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is read as 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol that provides information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements

must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this identifies a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to a separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? Specifies an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element (for example, 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! Specifies a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.
- * Specifies a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data-area, you know that you can include more than one data area or

you can include none. If you hear the lines 3*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.

+ Specifies a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times. For example, if you hear the line 6.1+ data-area, you must include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE, you know that you must include HOST, STATE, or both. As for the * symbol, you can only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. (enter the year or years).
All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

AIX; DB2; DB2 Universal Database; Distributed Relational Database Architecture; NUMA-Q; OS/2, OS/390, and OS/400; IBM Informix®; C-ISAM®; Foundation.2000™; IBM Informix® 4GL; IBM Informix® DataBlade® Module; Client SDK™; Cloudscape™; Cloudsync™; IBM Informix® Connect; IBM Informix® Driver for JDBC; Dynamic Connect™; IBM Informix® Dynamic Scalable Architecture™ (DSA); IBM Informix® Dynamic Server™; IBM Informix® Enterprise Gateway Manager (Enterprise Gateway Manager); IBM Informix® Extended Parallel Server™; i.Financial Services™; J/Foundation™; MaxConnect™; Object Translator™; Red Brick™; IBM Informix® SE; IBM Informix® SQL; InformiXML™; RedBack®; SystemBuilder™; U2™; UniData®; UniVerse®; wintegrate® are trademarks or registered trademarks of International Business Machines Corporation.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Windows, Windows NT, and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names used in this publication may be trademarks or service marks of others.

Index

A

- Accessibility xix
 - dotted decimal format of syntax diagrams A-1
 - syntax diagrams, reading in a screen reader A-1
- Administrative tool 4-17
- alter_java_path() 4-15

B

- Boldface type x

C

- Cast function, definition 4-2
- Circle class, example 6-5
- Class layout for input 6-8
- CLASS routine modifier 4-10
- CLASSPATH environment variable 2-4
- Code, sample, conventions for xv
 - com.informix.udr 4-4
 - com.informix.udr.UDREnv 4-5
 - com.informix.udr.UDRManager 4-5
 - com.informix.udr.UDRTraceable 4-7
 - com.informix.UDRLog 4-7
- Command-line conventions
 - how to read xiii
 - sample diagram xiii
- Compliance
 - with industry standards xxii
 - with SQLJ 4-19
- Configuration parameters
 - example 2-3
 - JDKVERSION 3-2
 - JVPCCLASSPATH 3-2
 - JVPHOME 3-3
 - JVPJAVAHOME 3-3
 - JVPJAVALIB 3-4
 - JVPJAVAVM 3-4
 - JVPLOGFILE 3-5
 - JVPPROFILE 3-6
 - SBSPACENAME 3-6
 - setting 2-3
 - VPCLASS JVP 3-7
- Configuring, to support Java 2-1
- Contact information xxiii
- Conventions
 - command-line xiii
 - documentation x
 - sample-code xv
 - syntax diagrams xii
 - syntax notation xii

- Conventions (*continued*)
 - typographical x
- COSTFUNC routine modifier 4-10

D

- Default I/O methods
 - and registering a UDT 6-2
 - backing out for opaque user-defined data type 6-2
 - for opaque user-defined data type 6-2
 - overriding 6-6
- Default locale ix
- Dependencies, software viii
- Deployment descriptor
 - example 4-13
 - in manifest file 4-14
 - SQL statements 4-13
- Disabilities, visual
 - reading syntax diagrams A-1
- Documentation conventions x
- Documentation Notes xvii
- Documentation set of all manuals xix
- Documentation, types of xvi
 - machine notes xvii
 - online manuals xix
 - printed manuals xix
- Dotted decimal format of syntax diagrams A-1
- DriverManager class 5-2

E

- en_us.8859-1 locale ix
- End-user routine 4-2
- Environment variables x
 - CLASSPATH 2-4
- Error messages xviii
- Example
 - circle class 6-5
 - circle UDT 4-9
 - configuration parameters 2-3
 - creating an sbospace 2-2
 - deployment descriptor 4-13
 - makefile for JAR 4-14
 - properties file 2-3
 - that overrides the default I/O methods 6-9
 - UDREnv class 4-5
 - usage of variable-length UDT 6-10
 - using CLASS 4-10
 - variable-length UDT 6-14
- EXECUTE FUNCTION statement 4-16
- EXECUTE PROCEDURE statement 4-16

F

Fixed and Known Defects File xvii
Functional index 4-3

G

Global Language Support (GLS) viii
GRANT statement 4-19

H

HANDLESNULLS routine modifier 4-10, 4-20
Help xix

I

I/O functions sets, for overriding default methods 6-6
IfmxIEBInStream 6-8
IfmxIEInStream 6-8
IfmxIEOutStream 6-8
IfmxSQLInStream 6-7
IfmxSQLOutStream 6-7
IfmxSRInStream 6-8
IfmxSROutStream 6-8
IfmxUDTSQLInput
 methods in interface 6-4
IfmxUDTSQLInput stream
 and IfxDatapointer class 6-7
IfmxUDTSQLOutput 6-3
 stream, and IfxDatapointer class 6-7
IfmxUDTSQLOutput interface
 methods contained in 6-5
IfxConnection 5-2
IfxDatapointer class 6-7
IfxDirectConnection 5-3
IfxDirectProtocol 5-3
IfxProtocol 5-2
Industry standards, compliance with xxii
Informix Dynamic Server documentation set xix
informix-direct subprotocol 5-2
informix.jvpcontrol 4-17
INFORMIXDIR/bin directory ix
install_jar() 4-15
Installation Guides xvi
Installing JDBC 2-1
Interface
 IfxConnection 5-2
 IfxProtocol 5-2
INTERNAL routine modifier 4-10
ISO 8859-1 code set ix
Iterator function 4-3
ITERATOR routine modifier 4-10
Iterator status 4-6

J

JAR file
 installing 4-15
 makefile 4-14

Java debugger 4-18
Java Development Kit 2-1
Java Runtime Environment viii, 2-1
Java Virtual Machine 1-1
Java virtual processor
 description 1-1
 log file 4-17
 messages 4-16
 specifying 4-9
java.sql.SQLInput
 extensions 5-5
 unsupported methods 5-4
java.sql.SQLOutput
 extensions 5-5
 unsupported methods 5-4
JDBC 1.0, extensions 5-4
JDBC 2.0
 Informix extensions 5-5
 public interfaces 5-5
JDKVERSION parameter 3-2
JVP properties file 2-2
JVP.monitor.port 4-17
JVPCLASSPATH parameter 3-2
JVPHOME parameter 3-3
JVPJAVAHOME parameter 3-3
JVPJAVALIB parameter 3-4
JVPJAVAVM parameter 3-4
JVPLOGFILE parameter 3-5
JVPPROFILE parameter 3-6

K

Keywords
 in syntax diagrams xiv

L

Limitations, streams 6-15
Locale viii
Log file 4-17
Logging 4-7

M

Machine notes xvii
Manifest file 4-14
Mapping
 between SQL and Java 4-8
 creating 4-9
Memory use 4-18

N

Negator function 4-3
NEGATOR routine modifier 4-10
NOT VARIANT routine modifier 4-10

O

Online help xix
Online manuals xix

Online notes xvi, xvii
onstat command 4-18
Opaque data type support function 4-3
Operator function 4-2
Overriding Default I/O Methods 6-6

P

Parallel queries 1-2
PARALLELIZABLE routine modifier 4-10
Parallelizable UDR 4-2
PERCALL_COST routine modifier 4-10
Port number 4-17
Printed manuals xix
Properties file 2-2

Q

Query parallelization 1-2

R

readsql() 6-2
registerDriver() 5-2
Registering a UDR 4-9
Registering a UDT
 default I/O methods 6-2
 registerJUDTfuncs 6-3
 with setUDTExName() SQL procedure 6-3
Release Notes xvii
remove_jar() 4-15
replace_jar() 4-15
Routine modifier
 CLASS 4-10
 COSTFUNC 4-10
 HANDLESNULLS 4-10
 INTERNAL 4-10
 ITERATOR 4-10
 NEGATOR 4-10
 NOT VARIANT 4-10
 PARALLELIZABLE 4-10
 PERCALL_COST 4-10
 SELCOST 4-10
 SELFUNC 4-10
 STACK 4-10
 unsupported 4-10, 4-20, 4-21
 VARIANT 4-10

S

Sample-code conventions xv
sbspace, creating 2-2
SBSPACENAME parameter 3-6
Screen reader
 reading syntax diagrams A-1
SELCOST routine modifier 4-10
SELECT statement 4-16
SELFUNC routine modifier 4-10
setUDTExtName 4-9
setUDTExtName(), registering a UDT 6-3

Software dependencies viii
SQL code xv
SQL definitions
 for default I/O UDRs 6-3
 for variable-length UDT example 6-14
SQL statement
 EXECUTE FUNCTION 4-16
 EXECUTE PROCEDURE 4-16
 GRANT 4-19
 SELECT 4-16
SQLBuffer
 transport type 6-3
SQLData interface
 readsql() 6-2
 using 6-1
 writesql() 6-2
SQLException 4-16
SQLInput interface
 extensions 6-3
 methods contained in 6-3
SQLOutput interface
 extensions 6-3
 methods contained in 6-4
SQLUDTInput 5-5
SQLUDTOutput 5-5
Stack dumps 4-16
STACK routine modifier 4-10
stores_demo database ix
Streams
 and database server I/O 6-1
 IfmxIEInStream 6-8
 IfmxIEOutStream 6-8
 IfmxSQLInStream 6-7
 IfmxSQLOutStream 6-7
 IfmxSRInStream 6-8
 IfmxSROutStream 6-8
 IfmxTextInStream 6-8
 IfmxTextOutStream 6-8
 limitations 6-15
 stream.available method 6-10
Syntax diagrams
 conventions for xii
 keywords in xiv
 reading in a screen reader A-1
 variables in xv
Syntax segment xiv
System catalog tables 1-2, 4-18
 syslangauth 4-19
 sysroutinesnags 4-18
System requirements
 database viii
 software viii
System trace events 4-18

T

- Threads
 - listing 4-18
 - scheduling 1-2
- TOC Notes xvii
- Trace outputs 4-16
- Trace-level settings 4-8
- Traceable events 4-18
- Typographical conventions x

U

- UDREnv 4-5
- UDRLog 4-7
- UDRManager 4-5
- UDRTraceable 4-7
- User-defined aggregate 4-2
- User-defined data type
 - converting to and from 6-1
 - implementing with default I/O methods 6-2
- User-defined function 4-2
- User-defined procedure 4-2
- User-defined routine
 - compiling 4-14
 - data type 4-6
 - definition of 4-1
 - executing 4-16
 - granting usage privilege 4-19
 - iterator status 4-6
 - log messages 4-17
 - logging 4-7
 - name 4-6
 - packages allowed 4-4
 - privileges 4-13, 4-19
 - registering 4-9
 - steps for creating 4-3
 - tracing and debugging 4-16
 - unsupported modifiers 4-20, 4-21
 - user-state pointer 4-6
 - uses not allowed 4-3
 - uses of 4-2
- Users, types of viii

V

- Variable-length UDT
 - example 6-10
- Variables, in syntax diagrams xv
- VARIANT routine modifier 4-10, 4-20
- Visual disabilities
 - reading syntax diagrams A-1
- VPCLASS JVP parameter 3-7
- vpid number 4-18

W

- writesql() 6-2

Z

- Zone-based tracing 4-7



Printed in USA

G251-2291-00

