



Java in Adaptive Server Enterprise

Adaptive Server® Enterprise

15.0.2

DOCUMENT ID: DC31652-01-1502-01

LAST REVISED: November 2008

Copyright © 2008 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

About This Book	ix
CHAPTER 1 An Introduction to Java in the Database	1
Advantages of Java in the database	1
Capabilities of Java in the database	2
Invoking Java methods in the database	2
Storing Java classes as datatypes	3
Storing and querying XML in the database	4
Standards	4
Java in the database: questions and answers	4
What are the key features?	5
How can I store Java instructions in the database?	5
How is Java executed in the database?	6
How can I use Java and SQL together?	6
What is the Java API?	7
How can I access the Java API from SQL?	7
Which Java classes are supported in the Java API?	8
Can I install my own Java classes?	8
Can I access data using Java?	8
Can I use the same classes on client and server?	8
How do I use Java classes in SQL?	9
Where can I find information about Java in the database?	9
What you cannot do with Java in the database	9
Sample Java classes	10
CHAPTER 2 Preparing for and Maintaining Java in the Database	11
The Java runtime environment	11
Java classes in the database	11
JDBC drivers	12
The Java VM	12
Configuring memory for Java in the database	13
Enabling the server for Java	13
Disabling the server for Java	13

Creating Java classes and JARs	14
Writing the Java code.....	14
Compiling Java code.....	14
Saving classes in a JAR file	15
Installing Java classes in the database.....	15
Using installjava	16
Referencing other Java-SQL classes.....	18
Viewing information about installed classes and JARs	18
Downloading installed classes and JARs.....	19
Removing classes and JARs	19
Retaining classes	20

CHAPTER 3	Using Java Classes in SQL.....	21
	General concepts	22
	Java considerations.....	22
	Java-SQL names.....	23
	Using Java classes as datatypes	23
	Creating and altering tables with Java-SQL columns.....	24
	Selecting, inserting, updating, and deleting Java objects.....	26
	Invoking Java methods in SQL	28
	Sample methods	29
	Exceptions in Java-SQL methods	29
	Representing Java instances	30
	Assignment properties of Java-SQL data items.....	31
	Datatype mapping between Java and SQL fields	33
	Character sets for data and identifiers	34
	Subtypes in Java-SQL data	34
	Widening conversions	35
	Narrowing conversions.....	35
	Runtime versus compile-time datatypes	36
	The treatment of nulls in Java-SQL data.....	36
	References to fields and methods of null instances	37
	Null values as arguments to Java-SQL methods	38
	Null values when using the SQL convert function	39
	Java-SQL string data	40
	Zero-length strings	40
	Type and void methods.....	41
	Java void instance methods	42
	Java void static methods	43
	Equality and ordering operations	44
	Evaluation order and Java method calls	45
	Columns	45
	Variables and parameters	46
	Static variables in Java-SQL classes	46

Java classes in multiple databases	47
Scope	48
Cross-database references	48
Inter-class transfers	49
Passing inter-class arguments	50
Temporary and work databases	50
Java classes	51

CHAPTER 4

Data Access Using JDBC 57

Overview	57
JDBC concepts and terminology	58
Differences between client- and server-side JDBC	58
Permissions	59
Using JDBC to access data	59
Overview of the JDBCExamples class	60
The main() and serverMain() methods	61
Obtaining a JDBC connection: the Connector() method	62
Routing the action to other methods: the doAction() method	63
Executing imperative SQL operations: the doSQL() method	63
Executing an update statement: the updater() method	63
Executing a select statement: the selector() method	64
Calling a SQL stored procedure: the caller() method	65
Error handling in the native JDBC driver	66
The JDBCExamples class	68
The main() method	69
The serverMain() method	69
The connector() method	70
The doAction() method	70
The doSQL() method	72
The updater() method	72
The selector() method	72
The caller() method	73

CHAPTER 5

SQLJ Functions and Stored Procedures 75

Overview	75
Compliance with SQLJ Part 1 specifications	76
General issues	76
Security and permissions	77
SQLJ Examples	77
Invoking Java methods in Adaptive Server	78
Using Sybase Central to manage SQLJ functions and procedures	80
SQLJ user-defined functions	81
Handling null argument values	84

Deleting a SQLJ function name.....	86
SQLJ stored procedures	86
Modifying SQL data.....	88
Using input and output parameters	90
Returning result sets	93
Viewing information about SQLJ functions and procedures	97
Advanced topics.....	97
Mapping Java and SQL datatypes	97
Using the command main method.....	101
SQLJ and Sybase implementation: a comparison	102
SQLJExamples class	105

CHAPTER 6 Debugging Java in the Database 109

Introduction to debugging Java	109
How the debugger works.....	109
Requirements for using the Java debugger	109
What you can do with the debugger.....	110
Using the debugger.....	110
Starting the debugger and connecting to the database.....	110
Compiling classes for debugging	111
Attaching to a Java VM	111
The Source window.....	112
Options	113
Setting breakpoints.....	114
Disconnecting from the database.....	116
A debugging tutorial	117
Before you begin	117
Start the Java debugger and connect to the database.....	117
Attach to a Java VM	118
Load source code into the debugger.....	118
Step through source code	119
Inspecting and modifying variables	120

CHAPTER 7 Network Access Using java.net..... 123

Overview	123
java.net classes.....	124
Setting up java.net	124
Example usage	125
Using socket classes.....	125
Using the URL class.....	128
User notes.....	130

CHAPTER 8	Reference Topics	131
	JDK requirement for Java classes in the server.....	131
	Assignments.....	132
	Assignment rules at compile-time	132
	Assignment rules at runtime	132
	Allowed conversions	133
	Transferring Java-SQL objects to clients	134
	Supported Java API packages, classes, and methods	134
	Supported Java packages and classes.....	135
	Unsupported Java packages, classes, and methods	135
	Unsupported java.sql methods and interfaces	136
	Invoking SQL from Java	137
	Special considerations	138
	Transact-SQL commands from Java methods.....	138
	Datatype mapping between Java and SQL.....	142
	Java-SQL identifiers	144
	Java-SQL class and package names.....	145
	Java-SQL column declarations	146
	Java-SQL variable declarations	147
	Java-SQL column references.....	147
	Java-SQL member references	148
	Java-SQL method calls	149
Glossary		153
Index		157

About This Book

Audience

This book is for Sybase System Administrators, Database Owners, and users who are familiar with the Java programming language and Transact-SQL®, the Sybase version of Structured Query Language (SQL). Familiarity with Java Database Connectivity (JDBC) is assumed for those who use these features.

How to use this book

This book will assist you in installing, configuring, and using Java classes and methods in the Adaptive Server database. It includes these chapters:

- Chapter 1, “An Introduction to Java in the Database,” provides an overview of Java in Adaptive Server, including a “questions and answers” section for both novice and experienced Java users.
- Chapter 2, “Preparing for and Maintaining Java in the Database,” describes the Java runtime environment and the steps for enabling Java on the server and installing Java classes.
- Chapter 3, “Using Java Classes in SQL,” describes how to use Java-SQL classes in your Adaptive Server database.
- Chapter 4, “Data Access Using JDBC,” describes how you use a JDBC driver (on the server or on the client) to perform SQL operations in Java.
- Chapter 5, “SQLJ Functions and Stored Procedures,” describes how you can enclose and use Java methods in SQL wrappers.
- Chapter 6, “Debugging Java in the Database,” describes how you use the Sybase debugger with Java.
- Chapter 7, “Network Access Using java.net,” describes how you can use java.net, a package that allows you to create networking applications over TCP/IP. It enables classes running in Adaptive Server to access different kinds of servers.
- Chapter 8, “Reference Topics,” provides information about datatype mapping, Java-SQL syntax, and other useful information.

In addition, a glossary provides descriptions of the Java and Java-SQL terms used in this book.

Note Information about XML in the SQL database, included in this book through version 12.5 of Adaptive Server, is now included in *XML Services in Adaptive Server Enterprise*.

Related documents

The Adaptive Server[®] Enterprise documentation set consists of the following:

- The release bulletin for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the release bulletin may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use the Sybase Technical Library.

- The *Installation Guide* for your platform – describes installation, upgrade, and configuration procedures for all Adaptive Server and related Sybase products.
- *What's New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server version 15.0, the system changes added to support those features, and changes that may affect your existing applications.
- *ASE Replicator User's Guide* – describes how to use the Adaptive Server Replicator feature of Adaptive Server to implement basic replication from a primary server to one or more remote Adaptive Servers.
- *Component Integration Services User's Guide* – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.
- The *Configuration Guide* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *Enhanced Full-Text Search Specialty Data Store User's Guide* – describes how to use the Full-Text Search feature with Verity to search Adaptive Server Enterprise data.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Historical Server User's Guide* – describes how to use Historical Server to obtain performance information for SQL Server[®] and Adaptive Server.

- *Java in Adaptive Server Enterprise* – describes how to install and use Java classes as datatypes, functions, and stored procedures in the Adaptive Server database.
- *Job Scheduler User's Guide* – provides instructions on how to install and configure, and create and schedule jobs on a local or remote Adaptive Server using the command line or a graphical user interface (GUI).
- *Messaging Service User's Guide* – describes how to use Real Time Messaging Services to integrate TIBCO Java Message Service and IBM WebSphere MQ messaging services with all Adaptive Server database applications.
- *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.
- *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
- *Performance and Tuning Series* – a series of books that explain how to tune Adaptive Server for maximum performance:
 - *Basics* – the basics for understanding and investigating performance questions in Adaptive Server.
 - *Locking and Concurrency Control* – describes how the various locking schemas can be used for improving performance in Adaptive Server, and how to select indexes to minimize concurrency.
 - *Query Processing and Abstract Plans* – describes how the optimizer processes queries and how abstract plans can be used to change some of the optimizer plans.
 - *Physical Database Tuning* – describes how to manage physical data placement, space allocated for data, and the temporary databases.
 - *Monitoring Adaptive Server with sp_sysmon* – describes how to monitor Adaptive Server's performance with sp_sysmon.
 - *Improving Performance with Statistical Analysis* – describes how Adaptive Server stores and displays statistics, and how to use the set statistics command to analyze server statistics.
 - *Using the Monitoring Tables* – describes how to query Adaptive Server's monitoring tables for statistical and diagnostic information.

-
- *Quick Reference Guide* – provides a comprehensive listing of the names and syntax for commands, functions, system procedures, extended system procedures, data types, and utilities in a pocket-sized book (regular size when viewed in PDF format).
 - *Reference Manual* – is a series of four books that contains the following detailed Transact-SQL information:
 - *Building Blocks* – Transact-SQL datatypes, functions, global variables, expressions, identifiers and wildcards, and reserved words.
 - *Commands* – Transact-SQL commands.
 - *Procedures* – Transact-SQL system procedures, catalog stored procedures, system extended stored procedures, and dbcc stored procedures.
 - *Tables* – Transact-SQL system tables and dbcc tables.
 - *System Administration Guide* –
 - *Volume 1* – provides an introduction to the basics of system administration, including a description of configuration parameters, resource issues, character sets, sort orders, and diagnosing system problems. The second part of this book is an in-depth description of security administration.
 - *Volume 2* – includes instructions and guidelines for managing physical resources, mirroring devices, configuring memory and data caches, managing multiprocessor servers and user databases, mounting and unmounting databases, creating and using segments, using the `reorg` command, and checking database consistency. The second half of this book describes how to back up and restore system and user databases.
 - *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Full-size available only in print version; a compact version is available in PDF format.
 - *Transact-SQL User's Guide* – documents Transact-SQL, the Sybase enhanced version of the relational database language. This manual serves as a textbook for beginning users of the database management system. This manual also contains descriptions of the pubs2 and pubs3 sample databases.
 - *Troubleshooting Series* (for release 15.0) –

- *Troubleshooting: Error Messages Advanced Resolutions* – contains troubleshooting procedures for problems that you may encounter when using Sybase® Adaptive Server® Enterprise. The problems addressed here are those which the Sybase Technical Support staff hear about most often
- *Troubleshooting and Error Messages Guide* – contains detailed instructions on how to resolve the most frequently occurring Adaptive Server error messages. Most of the messages presented here contain error numbers (from the master.sysmessages table), but some error messages do not have error numbers, and occur only in Adaptive Server's error log.
- *User Guide for Encrypted Columns* – describes how configure and use encrypted columns with Adaptive Server
- *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
- *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase Failover to configure an Adaptive Server as a companion server in a high availability system.
- *Unified Agent and Agent Management Console* – describes the Unified Agent, which provides runtime services to manage, monitor and control distributed Sybase resources.
- *Utility Guide* – documents the Adaptive Server utility programs, such as isql and bcp, which are executed at the operating system level.
- *Web Services User's Guide* – explains how to configure, use, and troubleshoot Web Services for Adaptive Server.
- *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using the Sybase DTM XA interface with X/Open XA transaction managers.
- *XML Services in Adaptive Server Enterprise* – describes the Sybase native XML processor and the Sybase Java-based XML support, introduces XML in the database, and documents the query and mapping functions that comprise XML Services.

Other sources of information

Use the Sybase Getting Started CD, the SyBooks CD, and the Sybase Product Manuals Web site to learn more about your product:

-
- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
 - The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

Some documentation may be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Certification Report.
- 3 In the Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Certification Report title to display the report.

❖ Finding the latest information on component certifications

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.
- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.

- 3 Select Search to display the availability and certification report for the selection.

❖ **Creating a personalized view of the Sybase Web site (including support pages)**

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

Sybase EBFs and software maintenance

❖ **Finding the latest information on EBFs and software maintenance**

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
- 3 Select a product.
- 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as a Technical Support Contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click Edit Roles to add the “Technical Support Contact” role to your MySybase profile.

- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Conventions

This section describes conventions used for Java and Transact_SQL in this book.

This book uses these font and syntax conventions for Java items:

- Classes, interfaces, methods, and packages are shown in Helvetica within paragraph text. For example:

`SybEventHandler` interface

`setBinaryStream()` method

com.Sybase.jdbx package

- Objects, instances, and parameter names are shown in italics. For example:

“In the following example, *ctx* is a *DirContext* object.”

“*eventHdler* is an instance of the *SybEventHandler* class that you implement.”

“The *classes* parameter is a string that lists specific classes you want to debug.”

- Java names are always case sensitive. For example, if a Java method name is shown as `Misc.stripLeadingBlanks()`, you must type the method name exactly as displayed.

Transact-SQL is a free-form language. There are no rules about the number of words you can put on a line or where you must break a line. However, for readability, all examples and most syntax statements in this manual are formatted so that each clause of a statement begins on a new line. Clauses that have more than one part extend to additional lines, which are indented. Complex commands are formatted using modified Backus Naur Form (BNF) notation.

Table 1 shows the conventions for syntax statements that appear in this manual:

Table 1: Font and syntax conventions for this manual

Element	Example
Command names, procedure names, utility names, and other keywords display in sans serif font.	<code>select</code> <code>sp_configure</code>
Database names and datatypes are in sans serif font.	<code>master database</code>
Book names, file names, variables, and path names are in italics.	<i>System Administration Guide</i> <i>sql.ini</i> file <i>column_name</i> \$SYBASE/ASE directory
Variables—or words that stand for values that you fill in—when they are part of a query or statement, are in italics in Courier font.	<code>select column_name</code> <code>from table_name</code> <code>where search_conditions</code>
Type parentheses as part of the command.	<code>compute row_aggregate (column_name)</code>
Double colon, equals sign indicates that the syntax is written in BNF notation. Do not type this symbol. Indicates “is defined as”.	<code>::=</code>
Curly braces mean that you must choose at least one of the enclosed options. Do not type the braces.	<code>{cash, check, credit}</code>

Element	Example
Brackets mean that to choose one or more of the enclosed options is optional. Do not type the brackets.	<code>[cash check credit]</code>
The comma means you may choose as many of the options shown as you want. Separate your choices with commas as part of the command.	<code>cash, check, credit</code>
The pipe or vertical bar () means you may select only one of the options shown.	<code>cash check credit</code>
An ellipsis (...) means that you can <i>repeat</i> the last unit as many times as you like.	<pre>buy thing = price [cash check credit] [, thing = price [cash check credit]]...</pre> <p>You must buy at least one thing and give its price. You may choose a method of payment: one of the items enclosed in square brackets. You may also choose to buy additional things: as many of them as you like. For each thing you buy, give its name, its price, and (optionally) a method of payment.</p>

- Syntax statements (displaying the syntax and all options for a command) appear as follows:

```
sp_dropdevice [device_name]
```

For a command with more options:

```
select column_name
from table_name
where search_conditions
```

In syntax statements, keywords (commands) are in normal font and identifiers are in lowercase. Italic font shows user-supplied words.

- Examples showing the use of Transact-SQL commands are printed like this:

```
select * from publishers
```

- Examples of output from the computer appear as follows:

```
pub_id      pub_name                city                state
-----
0736        New Age Books           Boston              MA
0877        Binnet & Hardley         Washington          DC
1389        Algodata Infosystems   Berkeley            CA

(3 rows affected)
```

Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Adaptive Server HTML documentation has been tested for compliance with U.S. government Section 508 Accessibility requirements. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

Note You might need to configure your accessibility tool for optimal use. Some screen readers pronounce text based on its case; for example, they pronounce ALL UPPERCASE TEXT as initials, and MixedCase Text as words. You might find it helpful to configure your tool to announce syntax conventions. Consult the documentation for your tool.

For information about how Sybase supports accessibility, see Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards.

If you need help

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.

An Introduction to Java in the Database

This chapter provides an overview of Java classes in Adaptive Server Enterprise.

Topic	Page
Advantages of Java in the database	1
Capabilities of Java in the database	2
Standards	4
Java in the database: questions and answers	4
Sample Java classes	10

Advantages of Java in the database

Adaptive Server provides a runtime environment for Java, which means that Java code can be executed in the server. Building a runtime environment for Java in the database server provides powerful new ways of managing and storing both data and logic.

- You can use the Java programming language as an integral part of Transact-SQL.
- You can reuse Java code in the different layers of your application—client, middle-tier, or server—and use them wherever makes most sense to you.
- Java in Adaptive Server provides a more powerful language than stored procedures for building logic into the database.
- Java classes become rich, user-defined data types.
- Methods of Java classes provide new functions accessible from SQL.

- Java can be used in the database without jeopardizing the integrity, security, and robustness of the database. Using Java does not alter the behavior of existing SQL statements or other aspects of non-Java relational database behavior.

Capabilities of Java in the database

Java in Adaptive Server allows you to:

- Invoke Java methods in the database
- Store Java classes as datatypes
- Store and query XML in the database

Invoking Java methods in the database

You can install Java classes in Adaptive Server, and then invoke the static methods of those classes in two ways:

- You can invoke the Java methods directly in SQL.
- You can wrap the methods in SQL names and invoke them as you would standard Transact-SQL stored procedures.

Invoking Java methods directly in SQL

The methods of an object-oriented language correspond to the functions of a procedural language. You can invoke methods stored in the database by referencing them, with name qualification, on instances for instance methods, and on either instances or classes for static (class) methods. You can invoke the method directly in, for example, Transact-SQL select lists and where clauses.

You can use static methods that return a value to the caller as user-defined functions (UDFs).

Certain restrictions apply when using Java methods in this way:

- If the Java method accesses the database through JDBC, result-set values are available only to the Java method, not to the client application.

- Output parameters are not supported. A method can manipulate the data it receives from a JDBC connection, but the only value it can return to its caller is a single return value declared as part of its definition.

Invoking Java methods as SQLJ stored procedures and functions

You can enclose Java static methods in SQL wrappers and use them exactly as you would Transact-SQL stored procedures or built-in functions. This functionality:

- Allows Java methods to return output parameters and result sets to the calling environment.
- Allows you to take advantage of traditional SQL syntax, metadata, and permission capabilities.
- Allows you to invoke SQLJ functions across databases.
- Allows you to use existing Java methods as SQLJ procedures and functions on the server, on the client, and on any SQLJ-compliant, third-party database.
- Complies with Part 1 of the standard specification. See “Standards” on page 4.

Storing Java classes as datatypes

With Java in the database, you can install pure Java classes in a SQL system, and then use those classes in a natural manner as datatypes in a SQL database. This capability adds a full object-oriented datatype extension mechanism to SQL, using a model that is widely understood and a language that is portable and widely available. The objects that you create and store with this facility are readily transferable to any Java-enabled environment, either in another SQL system or standalone Java environment.

This capability of using Java classes in the database has two different but complementary uses:

- It provides a type extension mechanism for SQL, which you can use for data that is created and processed in SQL.

- It provides a persistent data capability for Java, which you can use to store data in SQL that is created and processed (mainly) in Java. Java in Adaptive Server provides a distinct advantage over traditional SQL facilities: you do not need to map the Java objects into scalar SQL datatypes or store the Java objects as untyped binary strings.

Storing and querying XML in the database

Similar to Hypertext Markup Language (HTML), the eXtensible Markup Language (XML) allows you to define your own application-specific markup tags and is thus particularly suited for data interchange.

XML Services in Adaptive Server Enterprise describes the Sybase native XML processor and the Sybase Java-based XML support, introduces XML in the database, and documents the query and mapping functions that comprise XML Services.

Standards

The ANSI SQL standards specify SQL extensions for using Java facilities in SQL. The Java-SQL specifications are in the SQL standard, “Part 13: SQL Routines and Types Using the Java™ Programming Language (SQL/JRT).” This standard is referred to informally as “SQLJ.”

Sybase supports the SQLJ specifications for Java routines, and provides equivalent facilities for Java types. In addition, Sybase extends the standard. For example, Adaptive Server allows you to reference Java methods and classes directly in SQL.

Java in the database: questions and answers

Although this book assumes that readers are familiar with Java, there is much to learn about Java in a database. Sybase is not only extending the capabilities of the database with Java, but also extending the capabilities of Java with the database.

Both experienced and novice Java users should read this section. It uses a question-and-answer format to familiarize you with the basics of Java in Adaptive Server.

What are the key features?

All of these points are explained in detail in later sections. With Java in Adaptive Server, you can:

- Run Java in the database server using an internal Java Virtual Machine (Java VM).
- Call Java functions (methods) directly from SQL statements.
- Wrap Java methods in SQL aliases and call them as standard SQL stored procedures and built-in functions.
- Access SQL data from Java using an internal JDBC driver.
- Use Java classes as SQL datatypes.
- Save instances of Java classes in tables.
- Generate XML-formatted documents from raw data stored in Adaptive Server databases and, conversely, store XML documents and data extracted from them in Adaptive Server databases.
- Debug Java in the database.

How can I store Java instructions in the database?

Java is an object-oriented language. Its instructions (source code) come in the form of classes. You write and compile the Java instructions outside the database into compiled classes (byte code), which are binary files holding Java instructions.

You then install the compiled classes into the database, where they can be executed in the database server.

Adaptive Server is a runtime environment for Java classes. You need a Java development environment, such as Sybase PowerJ™ or Sun Microsystems Java Development Kit (JDK), to write and compile Java.

How is Java executed in the database?

To support Java in the database, Adaptive Server:

- Comes with its own Java VM, specifically developed for handling Java processing in the server.
- Uses its own JDBC driver that runs in the server and accesses a database.

The Sybase Java VM runs in the database environment. It interprets compiled Java instructions and runs them in the database server.

The Sybase Java VM meets the JCM specifications from Java Software; it is designed to work with the 2.0 version of the Java API. It supports public class and instance methods; classes inheriting from other classes; the Java API; and access to protected, public, and private fields. Some Java API functions that are not appropriate in a server environment, such as user interface elements, are not supported. All supported Java API packages and classes come with Adaptive Server.

The Adaptive Server Java VM is available at all times to perform a Java operation whenever it is required as part of the execution of a SQL statement. The database server starts the Java VM automatically when it is needed; you do not need to take any explicit action to start or stop the Java VM.

Client- and server-side JDBC

JDBC is the industry standard API for executing SQL in Java.

Adaptive Server provides a native JDBC driver. This driver is designed to maximize performance as it executes on the server because it does not need to communicate across the network. This driver permits Java classes installed in a database to use JDBC classes that execute SQL statements.

When JDBC classes are used within a client application, you typically must use *jConnect™* for JDBC™, the Sybase client-side JDBC database driver, to provide the classes necessary to establish a database connection.

How can I use Java and SQL together?

A guiding principle for the design of Java in the database is that it provides a natural, open extension to existing SQL functionality.

- *Java operations are invoked from SQL* – Sybase has extended the range of SQL expressions to include fields and methods of Java objects, so that you can include Java operations in a SQL statement.
- *Java methods as SQLJ stored procedures and functions* – you create a SQLJ alias for Java static methods, so that you can invoke them as standard SQL stored procedures and user-defined functions (UDFs).
- *Java classes become user-defined datatypes* – you store Java class instances using the same SQL statements as those used for traditional SQL datatypes.

You can use classes that are part of the Java API, and classes created and compiled by Java developers.

What is the Java API?

The Java Application Programming Interface (API) is a set of classes defined by Sun Microsystems. It provides a range of base functionality that can be used and extended by Java developers. It is the core of “what you can do” with Java.

The Java API offers considerable functionality in its own right. A large portion of the Java API is built in to any database that is enabled to use Java code—which includes the majority of nonvisual classes from the Java API already familiar to developers using the Sun Microsystems JDK.

How can I access the Java API from SQL?

You can use the Java API in stored procedures, in UDFs, and in SQL statements as extensions to the available built-in functions provided by SQL.

For example, the SQL function `PI(*)` returns the value for Pi. The Java API class `java.lang.Math` has a parallel field named `PI` that returns the same value. But `java.lang.Math` also has a field named `E` that returns the base of the natural logarithm, as well as a method that computes the remainder operation on two arguments as prescribed by the IEEE754 standard.

Which Java classes are supported in the Java API?

Not all Java API classes are supported in the database. Some classes, for example, the `java.awt` package that contains user interface components for applications, are not appropriate inside a database server. Other classes, including part of `java.io`, deal with writing information to a disk, and are also not supported in the database server environment. See Chapter 8, “Reference Topics,” for a list of supported and unsupported classes.

Can I install my own Java classes?

You can install your own Java classes into the database as, for example, a user-created `Employee` class or `Inventory` class that a developer designed, wrote, and compiled with a Java compiler.

User-defined Java classes can contain both information and methods. Once installed in a database, Adaptive Server lets you use these classes in all parts and operations of the database and execute their functionality (in the form of class or instance methods).

Can I access data using Java?

The JDBC interface is an industry standard designed to access database systems. The JDBC classes are designed to connect to a database, request data using SQL statements, and return results that can be processed in the client application.

You can connect from a client application to Adaptive Server Enterprise via JDBC, using `jConnect` or a JDBC/ODBC bridge. Adaptive Server also provides an internal JDBC driver, which permits Java classes installed in a database to use JDBC classes that execute SQL statements.

Can I use the same classes on client and server?

You can create Java classes that can be used on different levels of an enterprise application. You can integrate the same Java class into either the client application, a middle tier, or the database.

How do I use Java classes in SQL?

Using Java classes, whether user-defined or from the Java API, is a three-step activity:

- 1 Write or acquire a set of Java classes that you want to use as SQL datatypes, or as SQL aliases for static methods.
- 2 Install those classes in the Adaptive Server database.
- 3 Use those classes in SQL code:
 - Call class (static) methods of those classes as UDFs.
 - Declare the Java classes as datatypes of SQL columns, variables, and parameters. In this book, they are called Java-SQL columns, variables, and parameters.
 - Reference the Java-SQL columns, their fields, and their methods.
 - Wrap static methods in SQL aliases and use them as stored procedures or functions.

Where can I find information about Java in the database?

There are many books about Java and Java in the database. Two particularly useful books are:

- James Gosling, Bill Joy, Guy Steele, and Gilad Bracha, *The Java™ Language Specification, Second Edition*, Addison-Wesley, 2000.
- Seth White, Maydene Fisher, Rick Cattell, Graham Hamilton, and Mark Hapner, *JDBC™ API Tutorial and Reference*, Second Edition, Addison-Wesley, 1999.

What you cannot do with Java in the database

Adaptive Server is a runtime environment for Java classes, not a Java development environment.

You cannot perform these actions in the database:

- Edit class source files (*.java files).
- Compile Java class source files (*.java files).

- Execute Java APIs that are not supported, such as applet and visual classes.
- Use Java threading. Adaptive Server does not support `java.lang.Thread` and `java.lang.ThreadGroup`. If you attempt to spawn a thread, Adaptive Server throws `java.lang.UnsupportedOperationException`.
- Use the Java Native Interface (JNI).
- Use Java objects as parameters sent to a remote procedure call or received from a remote procedure call. They do not translate correctly.
- Sybase recommends that you do not use static variables in methods referenced by Java-SQL functions, SQLJ functions, or SQLJ stored procedures. The values returned for these variables may be unreliable as the scope of the static variable is implementation-dependent.

Sample Java classes

The chapters of this book use simple Java classes to illustrate basic principles for using Java in the database. You can find copies of these classes in the chapters that describe them and in the Sybase release directory in `$SYBASE/$SYBASE_ASE/sample/JavaXml/JavaXml.zip` (UNIX) or `%SYBASE%\Ase-15_0\sample\JavaXml\JavaXml.zip` (Windows NT).

Preparing for and Maintaining Java in the Database

This chapter describes the Java runtime environment, how to enable Java on the server, and how to install and maintain Java classes in the database.

Topic	Page
The Java runtime environment	11
Configuring memory for Java in the database	13
Enabling the server for Java	13
Creating Java classes and JARs	14
Installing Java classes in the database	15
Viewing information about installed classes and JARs	18
Downloading installed classes and JARs	19
Removing classes and JARs	19

The Java runtime environment

The Adaptive Server runtime environment for Java requires a Java VM, which is available as part of the database server, and the Sybase runtime Java classes, or Java API. If you are running Java applications on the client, you may also require the Sybase JDBC driver, jConnect, on the client.

Java classes in the database

You can use either of the following sources for Java classes:

- Sybase runtime Java classes
- User-defined classes

Sybase runtime Java classes

The Sybase Java VM supports a subset of JDK version 2.0 (UNIX and Windows NT) classes and packages.

The Sybase runtime Java classes are the low-level classes installed to Java-enable a database. They are downloaded automatically when Adaptive Server is installed and are available thereafter from `$SYBASE`
`/$SYBASE_ASE/lib/runtime.zip` (UNIX) or
`%SYBASE%\%SYBASE_ASE%\lib\runtime.zip` (Windows NT). You do not need to set the CLASSPATH environment variable specifically for Java in Adaptive Server.

Sybase does not support runtime Java packages and classes that assume a screen display, deal with networking and remote communications, or handle security. See Chapter 8, “Reference Topics” for a list of supported and unsupported packages and classes.

User-defined Java classes

You install user-defined classes into the database using the `installjava` utility. Once installed, these classes are available from other classes in the database and from SQL as user-defined datatypes.

JDBC drivers

The Sybase native JDBC driver that comes with Adaptive Server supports JDBC versions 1.1 and 1.2, and is compliant with several classes and methods of JDBC version 2.0. See Chapter 8, “Reference Topics,” for a complete list of supported and not supported classes and methods.

If your system requires a JDBC driver on the client, you must use `jConnect` version 5.5 or later, which supports JDBC version 2.0.

The Java VM

To ensure that each invoked method is executed as quickly as possible, Sybase provides a Java VM. The Java VM runs on the server. The Java VM requires little or no administration once installation is complete.

Configuring memory for Java in the database

Use the `sp_configure` system procedure to change memory allocations for Java in Adaptive Server. You can change the memory allocation for:

- size of global fixed heap – specifies memory space for internal data structures.
- size of process object fixed heap – specifies the total memory space available for all user connections using the Java VM.
- size of shared class heap – specifies the shared memory space for all Java classes called into the Java VM.

See “Java Services” in the *System Administration Guide* for complete information about these configuration parameters.

Enabling the server for Java

To enable the server and its databases for Java, enter this command from `isql`:

```
sp_configure "enable java", 1
```

Then shut down and restart the server.

By default, Adaptive Server is not enabled for Java. You cannot install Java classes or perform any Java operations until the server is enabled for Java.

You can increase or decrease the amount of memory available for Java in Adaptive Server and optimize performance using `sp_configure`. Java configuration parameters are described in the *System Administration Guide*.

Disabling the server for Java

To disable Java in the database, enter this command from `isql`:

```
sp_configure "enable java", 0
```

Creating Java classes and JARs

The Sybase-supported classes from the JDK are installed on your system when you install Adaptive Server version 12 or later. This section describes the steps for creating and installing your own Java classes.

To make your Java classes (or classes from other sources) available for use in the server, follow these steps:

- 1 Write and save the Java code that defines the classes.
- 2 Compile the Java code.
- 3 Create Java archive (JAR) files to organize and contain your classes.
- 4 Install the JARs/classes in the database.

Writing the Java code

Use the Sun Java SDK or a development tool such as Sybase PowerJ to write the Java code for your class declarations. Save the Java code in a file with an extension of *.java*. The name and case of the file must be the same as that of the class.

Note Make certain that any Java API classes used by your classes are among the supported API classes listed in Chapter 8, “Reference Topics”.

Compiling Java code

This step turns the class declaration containing Java code into a new, separate file containing bytecode. The name of the new file is the same as the Java code file but has an extension of *.class*. You can run a compiled Java class in a Java runtime environment regardless of the platform on which it was compiled or the operating system on which it runs.

Warning! Java classes that you install and use in the server must be compiled with JDK 1.2.2. If you compile a class with a later JDK, you will be able to install it in the server using the *installjava* utility, but you will get a *java.lang.ClassFormatError* exception when you attempt to use the class.

Saving classes in a JAR file

You can organize your Java classes by collecting related classes in packages and storing them in JAR files. JAR files allow you to install or remove related classes as a group.

Installing uncompressed JARs

To install Java classes in a database, save the classes or packages in a JAR file, in uncompressed form. To create an uncompressed JAR file that contains Java classes, use the Java `jar cf0` (“zero”) command.

In this UNIX example, the `jar` command creates an uncompressed JAR file that contains all `.class` files in the `jcsPackage` directory:

```
jar cf0 jcsPackage.jar jcsPackage/*.class
```

Installing compressed JARs

You can also install a compressed JAR file if you first expand the compressed file using the `x` option of the `jar` command. In this UNIX example, `abcPackage` is a compressed file.

- 1 Place the compressed JAR file in an empty directory and expand it:

```
jar xf0 abcPackage.jar
```

- 2 Delete the compressed JAR file so that it won't be included in the new, uncompressed JAR file:

```
rm abcPackage.jar
```

- 3 Create the uncompressed JAR file:

```
jar cf0 abcPackage.jar*
```

Installing Java classes in the database

To install Java classes from a client operating system file, use the `installjava` (UNIX) or `instjava` (Windows NT) utility from the command line.

See the *Adaptive Server Enterprise Utilities Guide* for detailed information about these utilities. Both utilities perform the same tasks; for simplicity, this document uses UNIX examples.

Using *installjava*

installjava copies a JAR file into the Adaptive Server system and makes the Java classes contained in the JAR available for use in the current database. The syntax is:

```
installjava
-f file_name
[-new | -update]
[-j jar_name]
[-S server_name ]
[-U user_name ]
[-P password ]
[-D database_name ]
[-I interfaces_file ]
[-a display_charset ]
[-J client_charset ]
[-z language ]
[-t timeout ]
```

For example, to install classes in the *addr.jar* file, enter:

```
installjava -f "/home/usera/jars/addr.jar"
```

The *-f* parameter specifies an operating system file that contains a JAR. You must use the complete path name for the JAR.

This section describes retained JAR files (using *-j*) and updating installed JARs and classes (using *new* and *update*). For more information about these and the other options available with *installjava*, see the *Utility Guide*.

Note When you install a JAR file, Application Server copies the file to a temporary table and then installs it from there. If you install a large JAR file, you may need to expand the size of *tempdb* using the *alter database* command.

Warning! Java classes that you install and use in the server must be compiled with JDK 1.2.2. If you compile a class with a later JDK, you will be able to install it in the server using the *installjava* utility, but you will get a *java.lang.ClassFormatError* exception when you attempt to use the class.

Retaining the JAR file

When a JAR is installed in a database, the server disassembles the JAR, extracts the classes, and stores them separately. The JAR is not stored in the database unless you specify *installjava* with the *-j* parameter.

Use of `-j` determines whether the Adaptive Server system retains the JAR specified in `installjava` or uses the JAR only to extract the classes to be installed.

- If you specify the `-j` parameter, Adaptive Server installs the classes contained in the JAR in the normal manner, and then retains the JAR and its association with the installed classes.
- If you do not specify the `-j` parameter, Adaptive Server does not retain any association of the classes with the JAR. This is the default option.

Sybase recommends that you specify a JAR name so that you can better manage your installed classes. If you retain the JAR file:

- You can remove the JAR and all classes associated with it, all at once, with the `remove java` statement. Otherwise, you must remove each class or package of classes one at a time.
- You can use `extractjava` to download the JAR to an operating system file. See “Downloading installed classes and JARs” on page 19.

Updating installed classes

The `new` and `update` clauses of `installjava` indicate whether you want new classes to replace currently installed classes.

- If you specify `new`, you cannot install a class with the same name as an existing class.
- If you specify `update`, you can install a class with the same name as an existing class, and the newly installed class replaces the existing class.

Warning! If you alter a class used as a column datatype by reinstalling a modified version of the class, make sure that the modified class can read and use existing objects (rows) in tables using that class as a datatype. Otherwise, you may be unable to access existing objects without reinstalling the original class.

Substitution of new classes for installed classes depends also on whether the classes being installed or the already installed classes are associated with a JAR. Thus:

- If you update a JAR, all classes in the existing JAR are deleted and replaced with classes in the new JAR.

- A class can be associated only with a single JAR. You cannot install a class in one JAR if a class of that same name is already installed and associated with another JAR. Similarly, you cannot install a class not-associated with a JAR if that class is currently installed and associated with a JAR.

You can, however, install a class in a retained JAR with the same name as an installed class not associated with a JAR. In this case, the class not associated with a JAR is deleted and the new class of the same name is associated with the new JAR.

If you want to reorganize your installed classes in new JARs, you may find it easier to first disassociate the affected classes from their JARs. See “Retaining classes” on page 20 for more information.

Referencing other Java-SQL classes

Installed classes can reference other classes in the same JAR file and classes previously installed in the same database, but they cannot reference classes in other databases.

If the classes in a JAR file do reference undefined classes, an error may result:

- If an undefined class is referenced directly in SQL, it causes a syntax error for “undefined class.”
- If an undefined class is referenced within a Java method that has been invoked, it throws a Java exception that may be caught in the invoked Java method or cause the general SQL exception described in “Exceptions in Java-SQL methods” on page 29.

The definition of a class can contain references to unsupported classes and methods as long as they are not actively referenced or invoked. Similarly, an installed class can contain a reference to a user-defined class that is not installed in the same database as long as the class is not instantiated or referenced.

Viewing information about installed classes and JARs

To view information about classes and JARs installed in the database, use `sp_helpjava`. The syntax is:

```
sp_helpjava ['class' [, name [, 'detail' | , 'depends' ]]] |  
            'jar' [, name [, 'depends' ] ]]
```

To view detailed information about the Address class, for example, log in to isql and enter:

```
sp_helpjava "class", Address, detail
```

See “sp_helpjava” in the *Reference Manual* for more information.

Downloading installed classes and JARs

You can download copies of Java classes installed on one database for use in other databases or applications.

Use the `extractjava` system utility to download a JAR file and its classes to a client operating system file. For example, to download `addr.jar` to `~/home/usera/jars/addrcopy.jar`, enter:

```
extractjava -j 'addr.jar' -f  
            '~/home/usera/jars/addrcopy.jar'
```

See the *Utility Guide* manual for more information.

Removing classes and JARs

Use the Transact-SQL `remove java` statement to uninstall one or more Java-SQL classes from the database. `remove java` can specify one or more Java class names, Java package names, or retained JAR names. For example, to uninstall the package `utilityClasses`, from isql enter:

```
remove java package "utilityClasses"
```

Note Adaptive Server does not allow you to remove classes that are used as the datatypes for columns and parameters or that are referenced by SQLJ functions or stored procedures.

You must make sure that you do not remove subclasses or classes that are used as variables or UDF return types.

`remove java package` deletes all classes in the specified package and all of its sub-packages.

See the *Reference Manual* for more information about `remove java`.

Retaining classes

You can delete a JAR file from the database but retain its classes as classes no longer associated with a JAR. Use `remove java` with the `retain classes` option if, for example, you want to rearrange the contents of several retained JARs.

For example, from `isql` enter:

```
remove java jar 'utilityClasses' retain classes
```

Once the classes are disassociated from their JARs, you can associate them with new JARs using `installjava` with the `new` keyword.

Using Java Classes in SQL

This chapter describes how to use Java classes in an Adaptive Server environment. The first sections give you enough information to get started; succeeding sections provide more advanced information.

Topics	Page
General concepts	22
Using Java classes as datatypes	23
Invoking Java methods in SQL	28
Representing Java instances	30
Assignment properties of Java-SQL data items	31
Datatype mapping between Java and SQL fields	33
Character sets for data and identifiers	34
Subtypes in Java-SQL data	34
The treatment of nulls in Java-SQL data	36
Java-SQL string data	40
Type and void methods	41
Equality and ordering operations	44
Evaluation order and Java method calls	45
Static variables in Java-SQL classes	46
Java classes in multiple databases	47
Java classes	51

In this document, SQL columns and variables whose datatypes are Java-SQL classes are described as Java-SQL columns and Java-SQL variables or as Java-SQL data items.

The sample classes used in this chapter can be found in:

- `$SYBASE/$SYBASE_ASE/sample/JavaXml/JavaXml.zip` (UNIX)
- `%SYBASE%\Ase-15_0\sample\JavaXml\JavaXml.zip` (Windows NT)

General concepts

This sections provides general Java and Java-SQL identifier information.

Java considerations

Before you use Java in your Adaptive Server database, here are some general considerations.

- Java classes contain:
 - Fields that have declared Java datatypes.
 - Methods whose parameters and results have declared Java datatypes.
 - Java datatypes for which there are corresponding SQL datatypes are defined in “Datatype mapping between Java and SQL” on page 142.
- Java classes can include classes, fields, and methods that are `private`, `protected`, `friendly`, or `public`.

Classes, fields and methods that are `public` can be referenced in SQL. Classes, fields, and methods that are `private`, `protected`, or `friendly` cannot be referenced in SQL, but they can be referenced in Java, and are subject to normal Java rules.

- Java classes, fields, and methods all have various syntactic properties:
 - Classes – the number of fields and their names
 - Field – their datatypes
 - Methods – the number of parameters and their datatypes, and the datatype of the result

The SQL system determines these syntactic properties from the Java-SQL classes themselves, using the Java Reflection API.

Java-SQL names

Java-SQL class names (identifiers) are limited to 255 bytes. Java-SQL field and method names can be any length, but they must be 255 bytes or less if you use them in Transact-SQL. All Java-SQL names must conform to the rules for Transact-SQL identifiers if you use them in Transact-SQL statements.

Class, field, and method names of 30 or more bytes must be surrounded by quotation marks.

The first character of the name must be either an alphabetic character (uppercase or lowercase) or an underscore (_) symbol. Subsequent characters can include alphabetic characters, numbers, the dollar (\$) symbol, or the underscore (_) symbol.

Java-SQL names are always case sensitive, regardless of whether the SQL system is specified as case sensitive or case insensitive.

See Java-SQL identifiers on page 144 for more information about identifiers.

Using Java classes as datatypes

After you have installed a set of Java classes, you can reference them as datatypes in SQL. To be used as a column datatype, a Java-SQL class must be defined as `public` and must implement either `java.io.Serializable` or `java.io.Externalizable`.

You can specify Java-SQL classes as:

- The datatypes of SQL columns
- The datatypes of Transact-SQL variables and parameters to Transact-SQL stored procedures
- Default values for SQL columns

When you create a table, you can specify Java-SQL classes as the datatypes of SQL columns:

```
create table emps (  
    name varchar(30),  
    home_addr Address,  
    mailing Address2Line null )
```

The `name` column is an ordinary SQL character string, the `home_addr` and `mailing_addr` columns can contain Java objects, and `Address` and `Address2Line` are Java-SQL classes that have been installed in the database.

You can specify Java-SQL classes as the datatypes of Transact-SQL variables:

```
declare @A Address
declare @A2 Address2Line
```

You can also specify default values for Java-SQL columns, subject to the normal constraint that the specified default must be a constant expression. This expression is normally a constructor invocation using the `new` operator with constant arguments, such as the following:

```
create table emps (
    name varchar(30),
    home_addr Address default new Address
        ('Not known', ''),
    mailing_addr Address2Line
)
```

Creating and altering tables with Java-SQL columns

When you create or alter tables with Java-SQL columns, you can specify any installed Java class as a column datatype. You can also specify how the information in the column is to be stored. Your choice of storage options affects the speed with which Adaptive Server references and updates the fields in these columns.

Column values for a row typically are stored “in-row,” that is, consecutively on the data pages allocated to a table. However, you can also store Java-SQL columns in a separate “off-row” location in the same way that text and image data items are stored. The default value for Java-SQL columns is off-row.

If a Java-SQL column is stored in-row:

- Objects stored in-row are processed more quickly than objects stored off-row.
- An object stored in-row can occupy up to approximately 16K bytes, depending on the page size of the database server and other variables. This includes its entire serialization, not just the values in its fields. A Java object whose runtime representation is more than the 16K limit generates an exception, and the command aborts.

If a Java-SQL column is stored off-row, the column is subject to the restrictions that apply to text and image columns:

- Objects stored off-row are processed more slowly than objects stored in-row.
- An object stored off-row can be of any size—subject to normal limits on text and image columns.
- An off-row column cannot be referenced in a check constraint.

Similarly, do not reference a table that contains an off-row column in a check constraint. Adaptive Server allows you to include the check constraint when you create or alter the table, but issues a warning message at compile time and ignores the constraint at runtime.

- You cannot include an off-row column in the column list of a select query with `select distinct`.
- You cannot specify an off-row column in a comparison operator, in a predicate, or in a `group by` clause.

Partial syntax for `create table` with the `in row/off row` option is:

```
create table...column_name datatype
    [default {constant_expression | user | null}]
    [{[identity | null | not null]}
    [off row | [ in row [ ( size_in_bytes ) ] ]]...
```

size_in_bytes specifies the maximum size of the in-row column. The value can be as large as 16K bytes. The default value is 255 bytes.

The maximum in-row column size you enter in `create table` must include the column's entire serialization, not just the values in its fields, plus minimum values for overhead.

To determine an appropriate column size that includes overhead and serialization values, use the `datalength` system function. `datalength` allows you to determine the actual size of a representative object you intend to store in the column.

For example:

```
select datalength (new class_name(...))
```

where *class_name* is an installed Java-SQL class.

Partial syntax for `alter table` is:

```
alter table...{add column_name datatype
               [default {constant_expression | user | null}]
               {identity | null} [ off row | [ in row ]...
```

Note You cannot change the column size of an in-row column using `alter column` in this Adaptive Server release.

Altering partitioned tables

If a table containing Java columns is partitioned, you cannot alter the table without first dropping the partitions. To change the table schema:

- 1 Remove the partitions.
- 2 Use the `alter table` command.
- 3 Repartition the table.

Selecting, inserting, updating, and deleting Java objects

After you specify Java-SQL columns, the values that you assign to those data items must be Java instances. Such instances are generated initially by calls to Java constructors using the `new` operator. You can generate Java instances for both columns and variables.

Constructor methods are pseudo instance methods. They create instances. Constructor methods have the same name as the class, and have no declared datatype. If you do not include a constructor method in your class definition, a default method is provided by the Java base class object. You can supply more than one constructor for each class, with different numbers and types of arguments. When a constructor is invoked, the one with the proper number and type of arguments is used.

In the following example, Java instances are generated for both columns and variables:

```
declare @A Address, @AA Address, @A2 Address2Line,
        @AA2 Address2Line

select @A = new Address( )
select @AA = new Address('123 Main Street', '99123')
select @A2 = new Address2Line( )
select @AA2 = new Address2Line('987 Front Street',
                               'Unit 2', '99543')
```

```
insert into emps values('John Doe', new Address( ),
    new Address2Line( ))
insert into emps values('Bob Smith',
    new Address('432 ElmStreet', '99654'),
    new Address2Line('PO Box 99', 'attn: Bob Smith', '99678') )
```

Values assigned to Java-SQL columns and variables can then be assigned to other Java-SQL columns and variables. For example:

```
declare @A Address, @AA Address, @A2 Address2Line,
    @AA2 Address2Line

select @A = home_addr, @A2 = mailing_addr from emps
    where name = 'John Doe'
insert into emps values ('George Baker', @A, @A2)

select @AA2 = @A2
update emps
    set home_addr = new Address('456 Shoreline Drive', '99321'),
        mailing_addr = @AA2
    where name = 'Bob Smith'
```

You can also copy values of Java-SQL columns from one table to another. For example:

```
create table trainees (
    name char(30),
    home_addr Address,
    mailing_addr Address2Line null
)
insert into trainees
select * from emps
    where name in ('Don Green', 'Bob Smith',
        'George Baker')
```

You can reference and update the fields of Java-SQL columns and of Java-SQL variables with normal SQL qualification. To avoid ambiguities with the SQL use of dots to qualify names, use a double-angle (>>) to qualify Java field and method names when referencing them in SQL.

```
declare @name varchar(100), @street varchar(100),
    @streetLine2 varchar(100), @zip char(10), @A Address

select @A = new Address()
select @A>>street = '789 Oak Lane'
select @street = @A>>street
```

```
select @street = home_addr>>street, @zip = home_addr>>zip from emps
      where name = 'Bob Smith'
select @name = name from emps
      where home_addr>>street= '456 Shoreline Drive'

update emps
      set home_addr>>street = '457 Shoreline Drive',
          home_addr>>zip = '99323'
      where home_addr>>street = '456 Shoreline Drive'
```

Invoking Java methods in SQL

You can invoke Java methods in SQL by referencing them, with name qualification, on instances for instance methods, and on either instances or classes for static methods.

Instance methods are generally closely tied to the data encapsulated in a particular instance of their class. Static (class) methods affect the whole class, not a particular instance of the class. Static methods often apply to objects and values from a wide range of classes.

Once you have installed a static method, it is ready for use. A class that contains a static method for use as a function must be `public`, but it does not need to be serializable.

One of the primary benefits of using Java with Adaptive Server is that you can use static methods that return a value to the caller as user-defined functions (UDFs).

You can use a Java static method as a UDF in a stored procedure, a trigger, a `where` clause, or anywhere that you can use a built-in SQL function.

Java methods invoked directly in SQL as UDFs are subject to these limitations:

- If the Java method accesses the database through JDBC, result-set values are available only to the Java method, not to the client application.
- Output parameters are not supported. A method can manipulate the data it receives from a JDBC connection, but the only value it can return to its caller is a single return value declared as part of its definition.
- Cross-database invocations of static methods are supported only if you use a class instance as a column value.

Permission to execute any UDF is granted implicitly to `public`. If the UDF performs SQL queries via JDBC, permission to access the data is checked against the invoker of the UDF. Thus, if user A invokes a UDF that accesses table `t1`, user A must have `select` permission on `t1` or the query will fail. For a more detailed discussion of security models for Java method invocations, see “Security and permissions” on page 77.

To use Java static methods to return result sets and output parameters, you must enclose the methods in SQL wrappers and invoke them as SQLJ stored procedures or functions. See “Invoking Java methods in Adaptive Server” on page 78 for a comparison of the ways you can invoke Java methods in Adaptive Server.

Sample methods

The sample `Address` and `Address2Line` classes have instance methods named `toString()`, and the sample `Misc` class has static methods named `stripLeadingBlanks()`, `getNumber()`, and `getStreet()`. You can invoke value methods as functions in a value expression.

```
declare @name varchar(100)
declare @street varchar(100)
declare @streetnum int
declare @A2 Address2Line

select @name = Misc.stripLeadingBlanks(name),
       @street = Misc.stripLeadingBlanks(home_addr>>street),
       @streetnum = Misc.getNumber(home_addr>>street),
       @A2 = mailing_addr
from emps
where home_addr>>toString( ) like '%Shoreline%'
```

For information about void methods (methods with no returned value) see “Type and void methods” on page 41.

Exceptions in Java-SQL methods

When the invocation of a Java-SQL method completes with unhandled exceptions, a SQL exception is raised, and this error message displays:

```
Unhandled Java method exception
```

The message text for the exception consists of the name of the Java class that raised the exception, followed by the character string (if any) supplied when the Java exception was thrown.

Representing Java instances

Non-Java clients such as `isql` cannot receive serialized Java objects from the server. To allow you to view and use the object, Adaptive Server must convert the object to a viewable representation.

To use an actual string value, Adaptive Server must invoke a method that translates the object into a `char` or `varchar` value. The `toString()` method in the `Address` class is an example of such a method. You must create your own version of the `toString()` method so that you can work with the viewable representation of the object.

Note The `toString()` method in the Java API does not convert the object to a viewable representation. The `toString()` method you create overrides the `toString()` method in the Java API.

When you use a `toString()` method, Adaptive Server imposes a limit on the number of bytes returned. Adaptive Server truncates the printable representation of the object to the value of the `@@stringsize` global variable. The default value of `@@stringsize` is 50; you can change this value using the `set stringsize` command. For example:

```
set stringsize 300
```

The display software on your computer may truncate the data item further so that it fits on the screen without wrapping.

If you include a `toString()` or similar method in each class, you can return the value of the object's `toString()` method in either of two ways:

- You can select a particular field in the Java-SQL column, which automatically invokes `toString()`:

```
select home__addr>>street from emps
```

- You can select the column and the `toString()` method, which lists in one string all of the field values in the column:

```
select home_addr>>toString() from emps
```


Assignment properties of Java-SQL data items

The values assigned to Java-SQL data items are derived ultimately from values constructed by Java-SQL methods in the Java VM. However, the logical representation of Java-SQL variables, parameters, and results is different from the logical representation of Java-SQL columns.

- Java-SQL *columns*, which are persistent, are Java serialized streams stored in the containing row of the table. They are stored values containing representations of Java instances.
- Java-SQL *variables*, *parameters*, and *function results* are transient. They do not actually contain Java-SQL instances, but instead contain references to Java instances contained in the Java VM.

These differences in representation give rise to differences in assignment properties as these examples illustrate.

- The `Address` constructor method with the `new` operator is evaluated in the Java VM. It constructs an `Address` instance and returns a reference to it. That reference is assigned as the value of Java-SQL variable `@A`:

```
declare @A Address, @AA Address, @A2 Address2Line,
        @AA2 Address2Line
select @A = new Address('432 Post Lane', '99444')
```

- Variable `@A` contains a reference to a Java instance in the Java VM. That reference is copied into variable `@AA`. Variables `@A` and `@AA` now reference the same instance.

```
select @AA=@A
```

- This assignment modifies the `zip` field of the `Address` referenced by `@A`. This is the same `Address` instance that is referenced by `@AA`. Therefore, the values of `@A.zip` and `@AA.zip` are now both '99222'.

```
select @A>>zip='99222'
```

- The `Address` constructor method with the `new` operator constructs an `Address` instance and returns a reference to it. However, since the target is a Java-SQL column, the SQL system serializes the `Address` instance denoted by that reference, and copies the serialized value into the new row of the `emps` table.

```
insert into emps
values ('Don Green', new Address('234 Stone
Road', '99777'), new Address2Line( ) )
```

The `Address2Line` constructor method operates the same way as the `Address` method, except that it returns a default instance rather than an instance with specified parameter values. The action taken is, however, the same as for the `Address` instance. The SQL system serializes the default `Address2Line` instance, and stores the serialized value into the new row of the `emps` table.

- The insert statement specifies no value for the `mailing_addr` column, so that column will be set to null, in the same manner as any other column whose value is not specified in an insert. This null value is generated entirely in SQL, and initialization of the `mailing_addr` column does not involve the Java VM at all.

```
insert into emps (name, home_addr) values ('Frank Lee', @A)
```

The insert statement specifies that the value of the `home_addr` column is to be taken from the Java-SQL variable `@A`. That variable contains a reference to an `Address` instance in the Java VM. Since the target is a Java-SQL column, the SQL system serializes the `Address` instance denoted by `@A`, and copies the serialized value into the new row of the `emps` table.

- This statement inserts a new `emps` row for 'Bob Brown.' The value of the `home_addr` column is taken from the SQL variable `@A`. It is also a serialization of the Java instance referenced by `@A`.

```
insert into emps (name, home_addr) values ('Bob Brown', @A)
```

- This update statement sets the `zip` field of the `home_addr` column of the 'Frank Lee' row to '99777.' This has no effect on the `zip` field in the 'Bob Brown' row, which is still '99444.'

```
update emps
  set home_add>>zip = '99777'
  where name = 'Frank Lee'
```

- The Java-SQL column `home_addr` contains a serialized representation of the value of an `Address` instance. The SQL system invokes the Java VM to deserialize that representation as a Java instance in the Java VM, and return a reference to the new deserialized copy. That reference is assigned to `@AA`. The deserialized `Address` instance that is referenced by `@AA` is entirely independent of both the column value and the instance referenced by `@A`.

```
select @AA = home_addr from emps where name = 'Frank Lee'
```

- This assignment modifies the zip field of the Address instance referenced by @A. This instance is a copy of the home_addr column of the 'Frank Lee' row, but is independent of that column value. The assignment therefore does not modify the zip field of the home_addr column of the 'Frank Lee' row.

```
select @A>>zip = '95678'
```

Datatype mapping between Java and SQL fields

When you transfer data in either direction between the Java VM and Adaptive Server, you must take into account that the datatypes of the data items are different in each system. Adaptive Server automatically maps SQL items to Java items and vice versa according to the correspondence tables in “Datatype mapping between Java and SQL” on page 142.

Thus, SQL type char translates to Java type String, the SQL type binary translates to the Java type byte[], and so on.

- For the datatype correspondences from SQL to Java, char, varchar, and varbinary types of any length correspond to Java String or byte[] datatypes, as appropriate.
- For the datatype correspondences from Java to SQL:
 - The Java String and byte[] datatypes correspond to SQL varchar and varbinary, where the maximum length value of 16K bytes is defined by Adaptive Server.
 - The Java BigDecimal datatype corresponds to SQL numeric(precision,scale), where precision and scale are defined by the user.

In the emps table, the maximum value for the Address and Address2Line classes, street, zip, and line2 fields is 255 bytes (the default value). The Java datatype of these classes is java.String, and they are treated in SQL as varchar(255).

An expression whose datatype is a Java object is converted to the corresponding SQL datatype only when the expression is used in a SQL context. For example, if the field home_addr>>street for employee 'Smith' is 260 characters, and begins '6789 Main Street ...:

```
select Misc.getStreet(home_addr>>street) from emps where name='Smith'
```

The expression in the select list passes the 260-character value of `home_addr>>street` to the `getStreet()` method (without truncating it to 255 characters). The `getStreet()` method then returns the 255-character string beginning 'Main Street....'. That 255-character string is now an element of the SQL select list, and is, therefore, converted to the SQL datatype and (if need be) truncated to 255 characters.

Character sets for data and identifiers

The character set for both Java source code and for Java String data is Unicode. Fields of Java-SQL classes can contain Unicode data.

Note Java identifiers used in the fully qualified names of visible classes or in the names of visible members can use only Latin characters and Arabic numerals.

Subtypes in Java-SQL data

Class subtypes allow you to use subtype substitution and method override, which are characteristics of Java. A conversion from a class to one of its superclasses is a widening conversion; a conversion from a class to one of its subclasses is a narrowing conversion.

- Widening conversions are performed implicitly with normal assignments and comparisons. They are always successful, since every subclass instance is also an instance of the superclass.
- Narrowing conversions must be specified with explicit `convert` expressions. A narrowing conversion is successful only if the superclass instance is an instance of the subclass, or a subclass of the subclass. Otherwise, an exception occurs.

Widening conversions

You do not need to use the `convert` function to specify a widening conversion. For example, since the `Address2Line` class is a subclass of the `Address` class, you can assign `Address2Line` values to `Address` data items. In the `emps` table, the `home_addr` column is an `Address` datatype and the `mailing_addr` column is an `Address2Line` datatype:

```
update emps
  set home_addr = mailing_addr
  where home_addr is null
```

For the rows fulfilling the `where` clause, the `home_addr` column contains an `Address2Line`, even though the declared type of `home_addr` is `Address`.

Such an assignment implicitly treats an instance of a class as an instance of a superclass of that class. The runtime instances of the subclass retain their subclass datatypes and associated data.

Narrowing conversions

You must use the `convert` function to convert an instance of a class to an instance of a subclass of the class. For example:

```
update emps
  set mailing_addr = convert(Address2Line, home_addr)
  where mailing_addr is null
```

The narrowing conversions in the `update` statement cause an exception if they are applied to any `home_addr` column that contains an `Address` instance that is not an `Address2Line`. You can avoid such exceptions by including a condition in the `where` clause:

```
update emps
  set mailing_addr = convert(Address2Line, home_addr)
  where mailing_addr is null
  and home_addr>>getClass( )>>toString( ) = 'Address2Line'
```

The expression “`home_addr>>getClass()>>toString()`” invokes `getClass()` and `toString()` methods of the Java `Object` class. The `Object` class is implicitly a superclass of all classes, so the methods defined for it are available for all classes.

You can also use a `case` expression:

```
update emps
  set mailing_addr =
```

```
case
  when home_addr>>getClass( )>>toString( )
    ='Address2Line'
  then convert(Address2Line, home_addr)
  else null
end
where mailing_addr is null
```

Runtime versus compile-time datatypes

Neither widening nor narrowing conversions modify the actual instance value or its runtime datatype; they simply specify the class to be used for the compile-time type. Thus, when you store `Address2Line` values from the `mailing_addr` column into the `home_address` column, those values still have the runtime type of `Address2Line`.

For example, the `Address` class and the `Address2Line` subclass both have the method `toString()`, which returns a `String` form of the complete address data.

```
select name, home_addr>>toString( ) from emps
where home_addr>>toString( ) not like '%Line2=[ ]'
```

For each row of `emps`, the declared type of the `home_addr` column is `Address`, but the runtime type of the `home_addr` value is either `Address` or `Address2Line`, depending on the effect of the previous update statement. For rows in which the runtime value of the `home_addr` column is an `Address`, the `toString()` method of the `Address` class is invoked, and for rows in which the runtime value of the `home_addr` column is `Address2Line`, the `toString()` method of the `Address2Line` subclass is invoked.

See “Null values when using the SQL `convert` function” on page 39 for a description of null values for widening and narrowing conversions.

The treatment of nulls in Java-SQL data

This section discusses the use of nulls in Java-SQL data items.

References to fields and methods of null instances

If the value of the instance specified in a field reference is null, then the field reference is null. Similarly, if the value of the instance specified in an instance method invocation is null, then the result of the invocation is null.

Java has different rules for the effect of referencing a field or method of a null instance. In Java, if you attempt to reference a field of a null instance, an exception is raised.

For example, suppose that the `emps` table has the following rows:

```
insert into emps (name, home_addr)
  values ("Al Adams",
    new Address("123 Main", "95321"))

insert into emps (name, home_addr)
  values ("Bob Baker",
    new Address("456 Side", "95123"))

insert into emps (name, home_addr)
  values ("Carl Carter", null)
```

Consider the following select:

```
select name, home_addr>>zip from emps
where home_addr>>zip in ('95123', '95125', '95128')
```

If the Java rule were used for the references to “`home_addr>>zip`,” then those references would cause an exception for the “Carl Carter” row, whose “`home_addr`” column is null. To avoid such an exception, you would need to write such a select as follows:

```
select name,
  case when home_addr is not null then home_addr>>zip
  else null end
from emps
  where case when home_addr is not null
  then home_addr>>zip
else
  null end
in ('95123', '95125', '95128')
```

The SQL convention is therefore used for references to fields and methods of null instances: if the instance is null, then any field or method reference is null. The effect of this SQL rule is to make the above case statement implicit.

However, this SQL rule for field references with null instances only applies to field references in source (right-side) contexts, not to field references that are targets (left-side) of assignments or set clauses. For example:

```
update emps
  set home_addr>>zip D '99123'
  where name D 'Charles Green'
```

This where clause is obviously true for the “Charles Green” row, so the update statement tries to perform the set clause. This raises an exception, because you cannot assign a value to a field of a null instance as the null instance has no field to which a value can be assigned. Thus, field references to fields of null instances are valid and return the null value in right-side contexts, and cause exceptions in left-side contexts.

The same considerations apply to invocations of methods of null instances, and the same rule is applied. For example, if we modify the previous example and invoke the `toString()` method of the `home_addr` column:

```
select name, home_addr>>toString( ) from emps
  where home_addr>>toString( ) D
    'StreetD234 Stone Road ZIPD 99777'
```

If the value of the instance specified in an instance method invocation is null, then the result of the invocation is null. Hence, the select statement is valid here, whereas it raises an exception in Java.

Null values as arguments to Java-SQL methods

The outcome of passing null as a parameter is independent of the actions of the method for which it is an argument, but instead depends on the ability of the return datatype to deliver a null value.

You cannot pass the null value as a parameter to a Java scalar type method; Java scalar types are always non-nullable. However, Java object types can accept null values.

For the following Java-SQL class:

```
public class General implements java.io.Serializable {
  public static int identity1(int I) {return I;}
  public static java.lang.Integer identity2
    (java.lang.Integer I) {return I;}
  public static Address identity3 (Address A) {return A;}
}
```

Consider these calls:


```

declare @I int
declare @A Address;

select @I = General.identity1(@I)
select @I = General.identity2(new java.lang.Integer(@I))
select @A = General.identity3(@A)

```

The values of both variable *@I* and variable *@A* are null, since values have not been assigned to them.

- The call of the `identity1()` method raises an exception. The datatype of the parameter *@I* of `identity1()` is the Java `int` type, which is scalar and has no null state. An attempt to pass a null valued argument to `identity1()` raises an exception.
- The call of the `identity2()` method succeeds. The datatype of the parameter of `identity2()` is the Java class `java.lang.Integer`, and the `new` expression creates an instance of `java.lang.Integer` that is set to the value of variable *@I*.
- The call of the `identity3()` method succeeds.

A successful call of `identity1()` never returns a null result because the return type has no null state. A null cannot be passed directly because the method resolution fails without parameter type information.

Successful calls of `identity2()` and `identity3()` can return null results.

Null values when using the SQL *convert* function

You use the `convert` function to convert a Java object of one class to a Java object of a superclass or subclass of that class.

As shown in “Subtypes in Java-SQL data” on page 34, the `home_addr` column of the `emps` table can contain values of both the `Address` class and the `Address2Line` class. In this example:

```

select name, home_addr>>street, convert(Address2Line, home_addr)>>line2,
       home_addr>>zip from emps

```

the expression “`convert(Address2Line, home_addr)`” contains a datatype (`Address2Line`) and an expression (`home_addr`). At compile-time, the expression (`home_addr`) must be a subtype or supertype of the class (`Address2Line`). At runtime, the action of this `convert` invocation depends on whether the runtime type of the expression’s value is a class, subclass, or superclass:

- If the runtime value of the expression (`home_addr`) is the specified class (`Address2Line`) or one of its subclasses, the value of the expression is returned, with the specified datatype (`Address2Line`).
- If the runtime value of the expression (`home_addr`) is a superclass of the specified class (`Address`), then a null is returned.

Adaptive Server evaluates the `select` statement for each row of the result. For each row:

- If the value of the `home_addr` column is an `Address2Line`, then `convert` returns that value, and the field reference extracts the `line2` field. If `convert` returns null, then the field reference itself is null.
- When a `convert` returns null, then the field reference itself evaluates to null.

Hence, the results of the `select` shows the `line2` value for those rows whose `home_addr` column is an `Address2Line` and a null for those rows whose `home_addr` column is an `Address`. As described in “The treatment of nulls in Java-SQL data” on page 36, the `select` also shows a null `line2` value for those rows in which the `home_addr` column is null.

Java-SQL string data

In Java-SQL columns, fields of type `String` are stored as Unicode.

When a Java-SQL `String` field is assigned to a SQL data item whose type is `char`, `varchar`, `nchar`, `nvarchar`, or `text`, the Unicode data is converted to the character set of the SQL system. Conversion errors are specified by the `set char_convert` options.

When a SQL data item whose type is `char`, `varchar`, `nchar`, or `text` is assigned to a Java-SQL `String` field that is stored as Unicode, the character data is converted to Unicode. Undefined codepoints in such data cause conversion errors.

Zero-length strings

In Transact-SQL, a zero-length character string is treated as a null value, and the empty string () is treated as a single space.

To be consistent with Transact-SQL, when a Java-SQL String value whose length is zero is assigned to a SQL data item whose type is char, varchar, nchar, nvarchar, or text, the Java-SQL String value is replaced with a single space.

For example:

```
1> declare @s varchar(20)
2> select @s = new java.lang.String()
3> select @s, char_length(@s)
4> go
```

(1 row affected)

```
-----
1
```

Otherwise, the zero-length value would be treated in SQL as a SQL null, and when assigned to a Java-SQL String, the Java-SQL String would be a Java null.

Type and void methods

Java methods (both instance and static) are either type methods or void methods. In general, type methods return a value with a result type, and void methods perform some action(s) and return nothing.

For example, in the Address class:

- The `toString()` method is a *type method* whose type is String.
- The `removeLeadingBlanks()` method is a *void method*.
- The Address constructor method is a *type method* whose type is the Address class.

You invoke type methods as functions and use the `new` keyword when invoking a constructor method:

```
insert into emps
  values ('Don Green', new Address('234 Stone Road', '99777'),
         new Address2Line( ) )

select name, home_addr>>toString( ) from emps
       where home_addr>>toString( ) like '%Baker%'
```

The `removeLeadingBlanks()` method of the `Address` class is a void instance method that modifies the `street` and `zip` fields of a given instance. You can invoke `removeLeadingBlanks()` for the `home_addr` column of each row of the `emps` table. For example:

```
update emps
  set home_addr =
    home_addr>>removeLeadingBlanks( )
```

`removeLeadingBlanks()` removes the leading blanks from the `street` and `zip` fields of the `home_addr` column. The Transact-SQL update statement does not provide a framework or syntax for such an action. It simply replaces column values.

Java void instance methods

To use the “update-in-place” actions of Java void instance methods in the SQL system, Java in Adaptive Server treats a call of a Java void instance method as follows:

For a void instance method `M()` of an instance `CI` of a class `C`, written “`CI.M(...)`”:

- In SQL, the call is treated as a type method call. The result type is implicitly class `C`, and the result value is a reference to `CI`. That reference identifies a copy of the instance `CI` after the actions of the void instance method call.
- In Java, this call is a void method call, which performs its actions and returns no value.

For example, you can invoke the `removeLeadingBlanks()` method for the `home_addr` column of selected rows of the `emps` table as follows:

```
update emps
  set home_addr = home_addr>>removeLeadingBlanks( )
where home_addr>>removeLeadingBlanks( )>>street like "123%"
```

- 1 In the `where` clause, “`home_addr>>removeLeadingBlanks()`” calls the `removeLeadingBlanks()` method for the `home_addr` column of a row of the `emps` table. `removeLeadingBlanks()` strips the leading blanks from the `street` and `zip` fields of a copy of the column. The SQL system then returns a reference to the modified copy of the `home_addr` column. The subsequent field reference:

```
home_addr>>removeLeadingBlanks( )>>street
```

returns the `street` field that has the leading blanks removed. The references to `home_addr` in the `where` clause are operating on a copy of the column. This evaluation of the `where` clause does *not* modify the `home_addr` column.

- 2 The `update` statement performs the `set` clause for each row of `emps` in which the `where` clause is true.
- 3 On the right-side of the `set` clause, the invocation of “`home_addr>>removeLeadingBlanks()`” is performed as it was for the `where` clause: `removeLeadingBlank()` strips the leading blanks from `street` and `zip` fields of that copy. The SQL system then returns a reference to the modified copy of the `home_addr` column.
- 4 The `Address` instance denoted by the result of the right side of the `set` clause is serialized and copied into the column specified on the left-side of the `set` clause: the result of the expression on the right side of the `set` clause is a copy of the `home_addr` column in which the leading blanks have been removed from the `street` and `zip` fields. The modified copy is then assigned back to the `home_addr` column as the new value of that column.

The expressions of the right and left side of the `set` clause are independent, as is normal for the `update` statement.

The following `update` statement shows an invocation of a void instance method of the `mailing_addr` column on the right side of the `set` clause being assigned to the `home_address` column on the left side.

```
update emps
  set home_addr = mailing_addr>>removeLeadingBlanks( )
  where ...
```

In this `set` clause, the void method `removeLeadingBlanks()` of the `mailing_addr` column yields a reference to a modified copy of the `Address2Line` instance in the `mailing_addr` column. The instance denoted by that reference is then serialized and assigned to the `home_addr` column. This action updates the `home_addr` column; it has no effect on the `mailing_addr` column.

Java void static methods

You cannot invoke a void static method using a simple SQL `execute` command. Rather, you must place the invocation of the void static method in a `select` statement.

For example, suppose that a Java class *C* has a void static method *M(...)*, and assume that *M()* performs an action you want to invoke in SQL. For example, *M()* can use JDBC calls to perform a series of SQL statements that have no return values, such as *create* or *drop*, that would be appropriate for a void method.

You must invoke the void static method in a *select* command, such as:

```
select C.M(...)
```

To allow void static methods to be invoked using a *select*, void static methods are treated in SQL as returning a value of datatype *int* with a value of null.

Equality and ordering operations

You can use equality and ordering operators when you use Java in the database. You cannot:

- Reference Java-SQL data items in ordering operations.
- Reference Java-SQL data items in equality operations if they are stored in an off-row column.
- Use the *order by* clause, which requires that you determine the sort order.
- Make direct comparisons using the “>”, “<”, “<=”, or “>=” operator.

These equality operations are allowed for in-row columns:

- Use of the *distinct* keyword, which is defined in terms of equality of rows, including Java-SQL columns.
- Direct comparisons using the “=” and “!=” operators.
- Use of the *union* operator (*not union all*), which eliminates duplicates, and requires the same kind of comparisons as the *distinct* clause.
- Use of the *group by* clause, which partitions the rows into sets with equal values of the grouping column.

Evaluation order and Java method calls

Adaptive Server does not have a defined order for evaluating operands of comparisons and other operations. Instead, Adaptive Server evaluates each query and chooses an evaluation order based on the most rapid rate of execution.

This section describes how different evaluation orders affect the outcome when you pass columns or variables and parameters as arguments. The examples in this section use the following Java-SQL class:

```
public class Utility implements java.io.Serializable {
    public static int F (Address A) {
        if (A.zip.length( ) > 5) return 0;
        else {A.zip = A.zip + "-1234"; return 1;}
    }
    public static int G (Address A) {
        if (A.zip.length( ) > 5) return 0;
        else {A.zip = A.zip + "-1234"; return 1;}
    }
}
```

Columns

In general, avoid invoking in the same SQL statement multiple methods on the same Java-SQL object. If at least one of them modifies the object, the order of evaluation can affect the outcome.

For example, in this example:

```
select * from emp E
where Utility.F(E.home_addr) > Utility.F(E.home_addr)
```

the where clause passes the same `home_addr` column in two different method invocations. Consider the evaluation of the where clause for a row whose `home_addr` column has a 5-character zip, such as “95123.”

Adaptive Server can initially evaluate either the left or right side of the comparison. After the first evaluation completes, the second is processed. Because it executes faster this way, Adaptive Server may let the second invocation see the modifications of the argument made by the first invocation.

In the example, the first invocation chosen by Adaptive Server returns 1, and the second returns 0. If the left operand is evaluated first, the comparison is $1 > 0$, and the `where` clause is true; if the right operand is evaluated first, the comparison is $0 > 1$, and the `where` clause is false.

Variables and parameters

Similarly, the order of evaluation can affect the outcome when passing variables and parameters as arguments.

Consider the following statements:

```
declare @A Address
declare @Order varchar(20)

select @A = new Address('95444', '123 Port Avenue')
select case when Utility.F(@A) > Utility.G(@A)
           then 'Left' else 'Right' end
select @Order = case when utility.F(@A) > utility.G(@A)
                   then 'Left' else 'Right' end
```

The new `Address` has a five-character zip code field. When the `case` expression is evaluated, depending on whether the left or right operand of the comparison is evaluated first, the comparison is either $1 > 0$ or $0 > 1$, and the `@Order` variable is set to 'Left' or 'Right' accordingly.

As for column arguments, the expression value depends on the evaluation order. Depending on whether the left or right operand of the comparison is evaluated first, the resulting value of the `zip` field of the `Address` instance referenced by `@A` is either "95444-4321" or "95444-1234."

Static variables in Java-SQL classes

A Java variable that is declared static is associated with the Java class, rather than with each instance of the class. The variable is allocated once for the entire class.

For example, you might include a static variable in the `Address` class that specifies the recommended limit on the length of the `Street` field:

```
public class Address implements java.io.Serializable {
    public static int recommendedLimit;
```



```
        public String street;  
        public String zip;  
    // ...  
}
```

You can specify that a static variable is final, which indicates that it is not updatable:

```
public static final int recommendedLimit;
```

Otherwise, you can update the variable.

You reference a static variable of a Java class in SQL by qualifying the static variable with an instance of the class. For example:

```
declare @a Address  
select @a>>recommendedLimit
```

If you don't have an instance of the class, you can use the following technique:

```
select convert(Address, null)>>recommendedLimit
```

The expression “(convert(null, Address))” converts a null value to an `Address` type; that is, it generates a null `Address` instance, which you can then qualify with the static variable name. You cannot reference a static variable of a Java class in SQL by qualifying the static variable with the class name. For example, the following are both incorrect:

```
select Address.recommendedLimit  
select Address>>recommendedLimit
```

Values assigned to non-final static variables are accessible only within the current session.

Java classes in multiple databases

You can store Java classes of the same name in different databases in the same Adaptive Server system. This section describes how you can use these classes.

Scope

When you install a Java class or set of classes, it is installed in the current database. When you dump or load a database, the Java-SQL classes that are currently installed in that database are always included—even if classes of the same name exist in other databases in the Adaptive Server system.

You can install Java classes with the same name in different databases. These synonymous classes can be:

- Identical classes that have been installed in different databases.
- Different classes that are intended to be mutually compatible. Thus, a serialized value generated by either class is acceptable to the other.
- Different classes that are intended to be “upward” compatible. That is, a serialized value generated by one of the classes should be acceptable to the other, but not vice versa.
- Different classes that are intended to be mutually incompatible; for example, a class named `Sheet` designed for supplies of paper, and other classes named `Sheet` designed for supplies of linen.

Cross-database references

You can reference objects stored in table columns in one database from another database.

For example, assume the following configuration:

- The `Address` class is installed in `db1` and `db2`.
- The `emps` table has been created in both `db1` with owner `Smith`, and in `db2`, with owner `Jones`.

In these examples, the current database is `db1`. You can invoke a join or a method across databases. For example:

- A join across databases might look like this:

```
declare @count int
select @count (*)
      from db2.Jones.emps, db1.Smith.emps
      where db2.Jones.emps.home_addr>>zip =
            db1.Smith.emps.home_addr>>zip
```

- A method invocation across databases might look like this:

```
select db2.Jones.emps.home_addr>>toString( )
```

```
from db2.Jones.emps
where db2.Jones.emps.name = 'John Stone'
```

In these examples, instance values are not transferred. Fields and methods of an instance contained in db2 are merely referenced by a routine in db1. Thus, for across-database joins and method invocations:

- db1 need not contain an Address class.
- If db1 does contain an Address class, it can have completely different properties than the Address class in db2.

Inter-class transfers

You can assign an instance of a class in one database to an instance of a class of the same name in another database. Instances created by the class in the source database are transferred into columns or variables whose declared type is the class in the current (target) database.

You can insert or update from a table in one database to a table in another database. For example:

```
insert into db1.Smith.emps select * from
    db2.Jones.emps

update db1.Smith.emps
    set home_addr = (select db2.Jones.emps.home_addr
                    from db2.Jones.emps
                    where db2.Jones.emps.name =
                        db1.Smith.emps.name)
```

You can insert or update from a variable in one database to another database. (The following fragment is in a stored procedure on db2.) For example:

```
declare @home_addr Address
select @home_addr = new Address('94608', '222 Baker
    Street')
insert into db1.Janes.emps(name, home_addr)
    values ('Jone Stone', @home_addr)
```

In these examples, instance values are transferred between databases. You can:

- Transfer instances between two local databases.
- Transfer instances between a local database and a remote database.
- Transfer instances between a SQL client and an Adaptive Server.

- Replace classes using `install` and `update` statements or `remove` and `update` statements.

In an inter-class transfer, the Java serialization is transferred from the source to the target. If the class in the source database is not compatible with the class in the target database, then the Java exception `InvalidClassException` is raised.

Passing inter-class arguments

You can pass arguments between classes of the same name in different databases. When passing inter-class arguments:

- A Java-SQL column is associated with the version of the specified Java class in the database that contains the column.
- A Java-SQL variable (in Transact-SQL) is associated with the version of the specified Java class in the current database.
- A Java-SQL intermediate result of class *C* is associated with the version of class *C* in the same database as the Java method that returned the result.
- When a Java instance value *JI* is assigned to a target variable or column, or passed to a Java method, *JI* is converted from its associated class to the class associated with the receiving target or method.

Temporary and work databases

All rules for Java classes and databases also apply to temporary databases and the model database:

- Java-SQL columns of temporary tables contain byte string serializations of the Java instances.
- A Java-SQL column is associated with the version of the specified class in the temporary database.

You can install Java classes in a temporary database, but they persist only as long as the temporary database persists.

The simplest way to provide Java classes for reference in temporary databases is to install Java classes in the model database. They are then present in any temporary database derived from the model.

Java classes

This section shows the simple Java classes that this chapter uses to illustrate Java in Adaptive Server. You can also find these classes and their Java source code in `$SYBASE/$SYBASE_ASE/sample/JavaXml/JavaXml.zip`. (UNIX) or `%SYBASE%\Ase-15_0\sample\JavaXml\JavaXml.zip` (Windows NT).

This is the Address class:

```
//
// Copyright (c) 2005
// Sybase, Inc
// Dublin, CA 94568
// All Rights Reserved
//
/**
 * A simple class for address data, to illustrate using a Java class
 * as a SQL datatype.
 */

public class Address implements java.io.Serializable {

    /**
     * The street data for the address.
     * @serial A simple String value.
     */
    public String street;

    /**
     * The zipcode data for the address.
     * @serial A simple String value.
     */
    String zip;

    /** A default constructor.
     */
    public Address ( ) {
        street = "Unknown";
        zip = "None";
    }

    /**
     * A constructor with parameters
     * @param S      a string with the street information
     * @param Z      a string with the zipcode information
     */
    public Address (String S, String Z) {
        street = S;
    }
}
```

```

        zip = Z;
    }

/**
 * A method to return a display of the address data.
 * @returns a string with a display version of the address data.
 */
    public String toString( ) {
        return "Street= " + street + " ZIP= " + zip;
    }

/**
 * A void method to remove leading blanks.
 * This method uses the static method
 * <code>Misc.stripLeadingBlanks</code>.
 */
    public void removeLeadingBlanks( ) {
        street = Misc.stripLeadingBlanks(street);
        zip = Misc.stripLeadingBlanks(zip);
    }
}

```

This is the Address2Line class, which is a subclass of the Address class:

```

//
// Copyright (c) 2005
// Sybase, Inc
// Dublin, CA 94568
// All Rights Reserved
//
/**
 * A subclass of the Address class that adds a second line of address data,
 * <p>This is a simple subclass to illustrate using a Java subclass
 * as a SQL datatype.
 */
public class Address2Line extends Address implements java.io.Serializable {

/**
 * The second line of street data for the address.
 * @serial a simple String value
 */
    String line2;

/**
 * A default constructor
 */
    public Address2Line ( ) {
        street = "Unknown";
        line2 = " ";
        zip = "None";
    }
}

```

```

    }

/**
 * A constructor with parameters.
 * @param S a string with the street information
 * @param L2 a string with the second line of address data
 * @param Z a string with the zipcode information
 */
public Address2Line (String S, String L2, String Z) {
    street = S;
    line2 = L2;
    zip = Z;
}

/**
 * A method to return a display of the address data
 * @returns a string with a display version of the address data
 */

public String toString( ) {
    return "Street= " + street + " Line2= " + line2 + " ZIP= " + zip;
}

/**
 * A void method to remove leading blanks.
 * This method uses the static method
 * <code>Misc.stripLeadingBlanks</code>.
 */

    public void removeLeadingBlanks( ) {
        line2 = Misc.stripLeadingBlanks(line2);
        super.removeLeadingBlanks( );
    }
}

```

The Misc class contains sets of miscellaneous routines:

```

//
// Copyright (c) 2005
// Sybase, Inc
// Dublin, CA 94568
// All Rights Reserved
//
/**
 * A non-instantiable class with miscellaneous static methods
 * that illustrate the use of Java methods in SQL.
 */

public class Misc{

```

```
/**
 * The Misc class contains only static methods and cannot be instantiated.
 */

private Misc( ) { }

/**
 * Removes leading blanks from a String
 */
    public static String stripLeadingBlanks(String s) {
        if (s == null) return null;
        for (int scan=0; scan<s.length( ); scan++)
            if (!java.lang.Character.isWhitespace(s.charAt(scan) ))
                break;
        } else if (scan == s.length( )){
            return "";
        } else return s.substring(scan);
        }
    }
}
return "";
}

/**
 * Extracts the street number from an address line.
 * e.g., Misc.getNumber(" 123 Main Street") == 123
 * Misc.getNumber(" Main Street") == 0
 * Misc.getNumber("") == 0
 * Misc.getNumber(" 123 ") == 123
 * Misc.getNumber(" Main 123 ") == 0
 * @param s a string assumed to have address data
 * @return a string with the extracted street number
 */

    public static int getNumber (String s) {
        String stripped = stripLeadingBlanks(s);
        if (s==null) return -1;
        for(int right=0; right < stripped.length( ); right++){
            if (!java.lang.Character.isDigit(stripped.charAt(right))) {
                break;
            } else if (right==0){
                return 0;
            } else {
                return java.lang.Integer.parseInt
                    (stripped.substring(0, right), 10);
            }
        }
    }
}
```



```
        }
        return -1;
    }

/**
 * Extract the "street" from an address line.
 * e.g., Misc.getStreet(" 123 Main Street") == "Main Street"
 *      Misc.getStreet(" Main Street") == "Main Street"
 *      Misc.getStreet("") == ""
 *      Misc.getStreet(" 123 ") == ""
 *      Misc.getStreet(" Main 123 ") == "Main 123"
 * @param s a string assumed to have address data
 * @return a string with the extracted street name
 */
    public static String getStreet(String s) {
        int left;
        if (s==null) return null;
        for (left=0; left<s.length( ); left++){
            if(java.lang.Character.isLetter(s.charAt(left))) {
                break;
            } else if (left == s.length( )) {
                return "";
            } else {
                return s.substring(left);
            }
        }
        return "";
    }
}
```


This chapter describes how to use Java Database Connectivity (JDBC) to access data.

Topics	Page
Overview	57
JDBC concepts and terminology	58
Differences between client- and server-side JDBC	58
Permissions	59
Using JDBC to access data	59
Error handling in the native JDBC driver	66
The JDBCExamples class	68

Overview

JDBC provides a SQL interface for Java applications. If you want to access relational data from Java, you must use JDBC calls.

You can use JDBC with the Adaptive Server SQL interface in either of two ways:

- *JDBC on the client* – Java client applications can make JDBC calls to Adaptive Server using the Sybase jConnect JDBC driver.
- *JDBC on the server* – Java classes installed in the database can make JDBC calls to the database using the JDBC driver native to Adaptive Server.

The use of JDBC calls to perform SQL operations is essentially the same in both contexts.

This chapter provides sample classes and methods that describe how you might perform SQL operations using JDBC. These classes and methods are not intended to serve as templates, but as general guidelines.

JDBC concepts and terminology

JDBC is a Java API and a standard part of the Java class libraries that control basic functions for Java application development. The SQL capabilities that JDBC provides are similar to those of ODBC and dynamic SQL.

The following sequence of events is typical of a JDBC application:

- 1 Create a *Connection* object – call the `getConnection()` static method of the `DriverManager` class to create a *Connection* object. This establishes a database connection.
- 2 Generate a *Statement* object – use the *Connection* object to generate a *Statement* object.
- 3 Pass a SQL statement to the *Statement* object – if the statement is a query, this action returns a *ResultSet* object.

The *ResultSet* object contains the data returned from the SQL statement, but provides it one row at a time (similar to the way a cursor works).

- 4 Loop over the rows of the results set – call the `next()` method of the *ResultSet* object to:
 - Advance the current row (the row in the result set that is being exposed through the *ResultSet* object) by one row.
 - Return a Boolean value (true/false) to indicate whether there is a row to advance to.
- 5 For each row, retrieve the values for columns in the *ResultSet* object – use the `getInt()`, `getString()`, or similar method to identify either the name or position of the column.

Differences between client- and server-side JDBC

The difference between JDBC on the client and in the database server is in how a connection is established with the database environment.

When you use client-side or server-side JDBC, you call the `DriverManager.getConnection()` method to establish a connection to the server.

- For client-side JDBC, you use the Sybase jConnect JDBC driver, and call the `DriverManager.getConnection()` method with the identification of the server. This establishes a connection to the designated server.

- For server-side JDBC, you use the Adaptive Server native JDBC driver, and call the `Drivermanager.getConnection()` method with one of the following values:
 - `jdbc:default:connection`
 - `jdbc:sybase:ase`
 - `jdbc:default`
 - empty string

This establishes a connection to the current server. Only the first call to the `getConnection()` method creates a new connection to the current server. Subsequent calls return a wrapper of that connection with all connection properties unchanged.

You can write JDBC classes to run at both the client and the server by using a conditional statement to set the URL.

Permissions

- *Java execution permissions* – like all Java classes in the database, classes containing JDBC statements can be accessed by any user. There is no equivalent of the `grant execute` statement that grants permission to execute procedures in Java methods, and there is no need to qualify the name of a class with the name of its owner.
- *SQL execution permissions* – Java classes are executed with the permissions of the connection executing them. This behavior is different from that of stored procedures, which execute with granted permission by the database owner.

Using JDBC to access data

This section describes how you can use JDBC to perform the typical operations of a SQL application. The examples are extracted from the class `JDBCExamples`, which is described in “The `JDBCExamples` class” on page 68 and in `$SYBASE/$SYBASE_ASE/sample/JavaXML/JavaXml.zip` (UNIX) or `%SYBASE%\Ase-12_5\sample\JavaXML\JavaXml.zip` (Windows NT).

JDBCExamples illustrates the basics of a user interface and shows the internal coding techniques for SQL operations.

Overview of the JDBCExamples class

The JDBCExamples class uses the Address class shown in “Sample Java classes” on page 10. To execute these examples on your machine, install the Address class on the server and include it in the Java CLASSPATH of the jConnect client.

You can call the methods of JDBCExamples from either a jConnect client or Adaptive Server.

Note You must create or drop stored procedures from the jConnect client. The Adaptive Server native driver does not support create procedure and drop procedure statements.

JDBCExamples static methods perform the following SQL operations:

- Create and drop an example table, xmp:

```
create table xmp (id int, name varchar(50), home Address)
```

- Create and drop a sample stored procedure, inoutproc:

```
create procedure inoutproc @id int, @newname varchar(50),  
    @newhome Address, @oldname varchar(50) output, @oldhome  
    Address output as
```

```
select @oldname = name, @oldhome = home from xmp  
where id=@id
```

```
update xmp set name=@newname, home = @newhome  
where id=@id
```

- Insert a row into the xmp table.
- Select a row from the xmp table.
- Update a row of the xmp table.
- Call the stored procedure inoutproc, which has both input parameters and output parameters of datatypes java.lang.String and Address.

JDBCExamples operates only on the xmp table and inoutproc procedure.

The *main()* and *serverMain()* methods

JDBCExamples has two primary methods:

- *main()* – is invoked from the command line of the jConnect client.
- *serverMain()* – performs the same actions as *main()*, but is invoked within Adaptive Server.

All actions of the JDBCExamples class are invoked by calling one of these methods, using a parameter to indicate the action to be performed.

Using *main()*

- You can invoke the *main()* method from a jConnect command line as follows:

```
java JDBCExamples
    "server-name:port-number?user=user-name&password=password" action
```

You can determine *server-name* and *port-number* from your interfaces file, using the dsedit tool. *user-name* and *password* are your user name and password. If you omit *&password=password*, the default is the empty password. Here are two examples:

```
"antibes:4000?user=smith&password=1x2x3"
"antibes:4000?user=sa"
```

Make sure that you enclose the parameter in quotation marks.

The *action* parameter can be create table, create procedure, insert, select, update, or call. It is case insensitive.

You can invoke JDBCExamples from a jConnect command line to create the table xmp and the stored procedure inoutproc as follows:

```
java JDBCExamples "antibes:4000?user=sa" CreateTable
java JDBCExamples "antibes:4000?user=sa" CreateProc
```

You can invoke JDBCExamples for insert, select, update, and call actions as follows:

```
java JDBCExamples "antibes:4000?user=sa" insert
java JDBCExamples "antibes:4000?user=sa" update
java JDBCExamples "antibes:4000?user=sa" call
java JDBCExamples "antibes:4000?user=sa" select
```

These invocations display the message “Action performed.”

To drop the table xmp and the stored procedure inoutproc, enter:

```
java JDBCExamples "antibes:4000?user=sa" droptable
java JDBCExamples "antibes:4000?user=sa" dropproc
```

Using `serverMain()`

Note Because the server-side JDBC driver does not support create procedure or drop procedure, create the table xmp and the example stored procedure inoutproc with client-side calls of the `main()` method before executing these examples. Refer to “Overview of the JDBCExamples class” on page 60.

After creating xmp and inoutproc, you can invoke the `serverMain()` method as follows:

```
select JDBCExamples.serverMain('insert')
go
select JDBCExamples.serverMain('select')
go
select JDBCExamples.serverMain('update')
go
select JDBCExamples.serverMain('call')
go
```

Note Server-side calls of `serverMain()` do not require a *server-name:port-number* parameter; Adaptive Server simply connects to itself.

Obtaining a JDBC connection: the `Connector()` method

Both `main()` and `serverMain()` call the `connector()` method, which returns a JDBC *Connection* object. The *Connection* object is the basis for all subsequent SQL operations.

Both `main()` and `serverMain()` call `connector()` with a parameter that specifies the JDBC driver for the server- or client-side environment. The returned *Connection* object is then passed as an argument to the other methods of the JDBCExamples class. By isolating the connection actions in the `connector()` method, JDBCExamples’ other methods are independent of their server- or client-side environment.

Routing the action to other methods: the *doAction()* method

The *doAction()* method routes the call to one of the other methods, based on the *action* parameter.

doAction() has the *Connection* parameter, which it simply relays to the target method. It also has a parameter *locale*, which indicates whether the call is server- or client-side. *Connection* raises an exception if either *create procedure* or *drop procedure* is invoked in a server-side environment.

Executing imperative SQL operations: the *doSQL()* method

The *doSQL()* method performs SQL actions that require no input or output parameters such as *create table*, *create procedure*, *drop table*, and *drop procedure*.

doSQL() has two parameters: the *Connection* object and the SQL statement it is to perform. *doSQL()* creates a JDBC *Statement* object and uses it to execute the specified SQL statement.

Executing an *update* statement: the *updater()* method

The *updater()* method performs a Transact-SQL *update* statement. The *update* action is:

```
String sql = "update xmp set name = ?, home = ? where id = ?";
```

It updates the *name* and *home* columns for all rows with a given *id* value.

The *update* values for the *name* and *home* column, and the *id* value, are specified by parameter markers (?). *updater()* supplies values for these parameter markers after preparing the statement, but before executing it. The values are specified by the JDBC *setString()*, *setObject()*, and *setInt()* methods with these parameters:

- The ordinal parameter marker to be substituted
- The value to be substituted

For example:

```
pstmt.setString(1, name);  
pstmt.setObject(2, home);  
pstmt.setInt(3, id);
```

After making these substitutions, *updater()* executes the *update* statement.

To simplify `updater()`, the substituted values in the example are fixed. Normally, applications compute the substituted values or obtain them as parameters.

Executing a *select* statement: the *selector()* method

The `selector()` method executes a Transact-SQL *select* statement:

```
String sql = "select name, home from xmp where id=?";
```

The *where* clause uses a parameter marker (?) for the row to be selected. Using the JDBC `setInt()` method, `selector()` supplies a value for the parameter marker after preparing the SQL statement:

```
PreparedStatement pstmt =  
    con.prepareStatement(sql);  
pstmt.setInt(1, id);
```

`selector()` then executes the *select* statement:

```
ResultSet rs = pstmt.executeQuery();
```

Note For SQL statements that return no results, use `doSQL()` and `updater()`. They execute SQL statements with the `executeUpdate()` method.

For SQL statements that do return results, use the `executeQuery()` method, which returns a JDBC *ResultSet* object.

The *ResultSet* object is similar to a SQL cursor. Initially, it is positioned before the first row of results. Each call of the `next()` method advances the *ResultSet* object to the next row, until there are no more rows.

`selector()` requires that the *ResultSet* object have exactly one row. The `selector()` method invokes the `next` method, and checks for the case where *ResultSet* has no rows or more than one row.

```
if (rs.next()) {  
    name = rs.getString(1);  
    home = (Address)rs.getObject(2);  
    if (rs.next()) {  
        throw new Exception("Error:  Select returned multiple rows");  
    } else { // No action  
    }  
} else { throw new Exception("Error:  Select returned no rows");  
}
```

In the above code, the call of methods `getString()` and `getObject()` retrieve the two columns of the first row of the result set. The expression `“(Address)rs.getObject(2)”` retrieves the second column as a Java object, and then coerces that object to the `Address` class. If the returned object is not an `Address`, then an exception is raised.

`selecter()` retrieves a single row and checks for the cases of no rows or more than one row. An application that processes a multiple row *ResultSet* would simply loop on the calls of the `next()` method, and process each row as for a single row.

Executing in batch mode

If you want to execute a batch of SQL statements, make sure that you use the `execute()` method. If you use `executeQuery()` for batch mode:

- If the batch operation does not return a result set (contains no `select` statements), the batch executes without error.
- If the batch operation returns one result set, all statements after the statement that returns the result are ignored. If `getXXX()` is called to get an output parameter, the remaining statements execute and the current result set is closed.
- If the batch operation returns more than one result set, an exception is raised and the operation aborts.

Using `execute()` ensures that the complete batch executes for all cases.

Calling a SQL stored procedure: the *caller()* method

The `caller()` method calls the stored procedure `inoutproc`:

```
create proc inoutproc @id int, @newname varchar(50), @newhome Address,
    @oldname varchar(50) output, @oldhome Address output as

select @oldname = name, @oldhome = home from xmp where id=@id
update xmp set name=@newname, home = @newhome where id=@id
```

This procedure has three input parameters (`@id`, `@newname`, and `@newhome`) and two output parameters (`@oldname` and `@oldhome`). `caller()` sets the name and home columns of the row of table `xmp` with the ID value of `@id` to the values `@newname` and `@newhome`, and returns the former values of those columns in the output parameters `@oldname` and `@oldhome`.

The `inoutproc` procedure illustrates how to supply input and output parameters in a JDBC call.

`caller()` executes the following call statement, which prepares the call statement:

```
CallableStatement cs = con.prepareCall("{call inoutproc (?, ?, ?, ?, ?)}");
```

All of the parameters of the call are specified as parameter markers (?).

`caller()` supplies values for the input parameters using JDBC `setInt()`, `setString()`, and `setObject()` methods that were used in the `doSQL()`, `updateAction()`, and `selector()` methods:

```
cs.setInt(1, id);
cs.setString(2, newName);
cs.setObject(3, newHome);
```

These set methods are not suitable for the output parameters. Before executing the call statement, `caller()` specifies the datatypes expected of the output parameters using the JDBC `registerOutParameter()` method:

```
cs.registerOutParameter(4, java.sql.Types.VARCHAR);
cs.registerOutParameter(5, java.sql.Types.JAVA_OBJECT);
```

`caller()` then executes the call statement and obtains the output values using the same `getString()` and `getObject()` methods that the `selector()` method used:

```
int res = cs.executeUpdate();
String oldName = cs.getString(4);
Address oldHome = (Address)cs.getObject(5);
```

Error handling in the native JDBC driver

Sybase supports and implements all methods from the `java.sql.SQLException` and `java.sql.SQLWarning` classes. `SQLException` provides information on database access errors. `SQLWarning` extends `SQLException` and provides information on database access warnings.

Errors raised by Adaptive Server are numbered according to severity. Lower numbers are less severe; higher numbers are more severe. Errors are grouped according to severity:

- Warnings (EX_INFO: severity 10) – are converted to `SQLWarnings`.
- Exceptions (severity 11 to 18) – are converted to `SQLExceptions`.
- Fatal errors (severity 19 to 24) – are converted to fatal `SQLExceptions`.

SQLExceptions can be raised through JDBC, Adaptive Server, or the native JDBC driver. Raising a SQLException aborts the JDBC query that caused the error. Subsequent system behavior differs depending on where the error is caught:

- *If the error is caught in Java* – a “try” block and subsequent “catch” block process the error.

Adaptive Server provides several extended JDBC driver-specific SQLException error messages. All are EX_USER (severity 16) and can always be caught. There are no driver-specific SQLWarning messages.

- *If the error is not caught in Java* – the Java VM returns control to Adaptive Server, Adaptive Server catches the error, and an unhandled SQLException error is raised.

The raiserror command is used typically with stored procedures to raise an error and to print a user-defined error message. When a stored procedure that calls the raiserror command is executed via JDBC, the error is treated as an internal error of severity EX_USER, and a nonfatal SQLException is raised.

Note You cannot access extended error data using the raiserror command; the with errordata clause is not implemented for SQLException.

If an error causes a transaction to abort, the outcome depends on the transaction context in which the Java method is invoked:

- *If the transaction contains multiple statements* – the transaction aborts and control returns to the server, which rolls back the entire transaction. The JDBC driver ceases to process queries until control returns from the server.
- *If the transaction contains a single statement* – the transaction aborts, the SQL statement it contains rolls back, and the JDBC driver continues to process queries.

The following scenarios illustrate the different outcomes. Consider a Java method jdbcTests.Errorexample() that contains these statements:

```
stmt.executeUpdate("delete from parts where partno = 0");           Q2
stmt.executeQuery("select 1/0");                                   Q3
stmt.executeUpdate("delete from parts where partno = 10");         Q4
```

A transaction containing multiple statements includes these SQL commands:

```
begin transaction
delete from parts where partno = 8                                Q1
select JDBCTests.Errorexample()
```

In this case, these actions result from an aborted transaction:

- A divide-by-zero exception is raised in Q3.
- Changes from Q1 and Q2 are rolled back.
- The entire transaction aborts.

A transaction containing a single statement includes these SQL commands:

```
set chained off
delete from parts where partno = 8           Q1
select JDBCTests.Errorexample()
```

In this case:

- A divide-by-zero exception is raised in Q3.
- Changes from Q1 and Q2 are not rolled back
- The exception is caught in “catch” and “try” blocks in JDBCTests.Errorexample.
- The deletion specified in Q4 does not execute because it is handled in the same “try” and “catch” blocks as Q3.
- JDBC queries outside of the current “try” and “catch” blocks can be executed.

The JDBCExamples class

```
// An example class illustrating the use of JDBC facilities
// with the Java in Adaptive Server feature.
//
// The methods of this class perform a range of SQL operations.
// These methods can be invoked either from a Java client,
// using the main method, or from the SQL server, using
// the serverMain method.
//
import java.sql.*;           // JDBC
public class JDBCExamples {
{
```

The main() method

```
// The main method, to be called from a client-side command line
//
public static void main(String args[]) {
    if (args.length!=2) {
        System.out.println("\n Usage:      "
            + "java ExternalConnect server-name:port-number
            action ");
        System.out.println(" The action is connect, createtable,
            " + "createproc, drop, "
            + "insert, select, update, or call \n" );
        return;
    }
    try{
        String server = args[0];
        String action = args[1].toLowerCase();
        Connection con = connector(server);
        String workString = doAction( action, con, client);
        System.out.println("\n" + workString + "\n");
    } catch (Exception e) {
        System.out.println("\n Exception: ");
        e.printStackTrace();
    }
}
```

The serverMain() method

```
// A JDBCExamples method equivalent to 'main',
// to be called from SQL or Java in the server

public static String serverMain(String action) {
    try {
        Connection con = connector("default");
        String workString = doAction(action, con, server);
        return workString;
    } catch ( Exception e ) {
        if (e.getMessage().equals(null)) {
            return "Exc: " + e.toString();
        } else {
            return "Exc - " + e.getMessage();
        }
    }
}
```

The connector() method

```
// A JDBCExamples method to get a connection.
// It can be called from the server with argument 'default',
// or from a client, with an argument that is the server name.

public static Connection connector(String server)
    throws Exception, SQLException, ClassNotFoundException {

    String forName="";
    String url="";

    if (server=="default") { // server connection to current server
        forName = "sybase.asejdbc.ASEDriver";
        url = "jdbc:default:connection";
    } else if (server!="default") { //client connection to server
        forName= "com.sybase.jdbc.SybDriver";
        url = "jdbc:sybase:Tds:"+ server;
    }

    String user = "sa";
    String password = "";

    // Load the driver
    Class.forName(forName);
    // Get a connection
    Connection con = DriverManager.getConnection(url,
        user, password);
    return con;
}
```

The doAction() method

```
// A JDBCExamples method to route to the 'action' to be performed

public static String doAction(String action, Connection con,
    String locale)
    throws Exception {

    String createProcScript =
        " create proc inoutproc @id int, @newname varchar(50),
        @newhome Address, "
        + "    @oldname varchar(50) output, @oldhome Address
        output as "
        + " select @oldname = name, @oldhome = home from xmp
```



```

        where id=@id "
    + " update xmp set name=@newname, home = @newhome
        where id=@id ";
String createTableScript =
    " create table xmp (id int, name varchar(50),
        home Address) " ;

String dropTableScript = "drop table xmp ";
String dropProcScript = "drop proc inoutproc ";
String insertScript = "insert into xmp "
    + "values (1, 'Joe Smith', new Address('987 Shore',
        '12345'))";

String workString = "Action (" + action + ) ;
if (action.equals("connect")) {
    workString += "performed";
} else if (action.equals("createtable")) {
    workString += doSQL(con, createTableScript );
} else if (action.equals("createproc")) {
    if (locale.equals(server)) {
        throw new exception (CreateProc cannot be performed
            in the server);
    } else {
        workString += doSQL(con, createProcScript );
    }
} else if (action.equals("droptable")) {
    workString += doSQL(con, dropTableScript );
} else if (action.equals("dropproc")) {
    if (locale.equals(server)) {
        throw new exception (CreateProc cannot be performed
            in the server);
    } else {
        workString += doSQL(con, dropProcScript );
    }
} else if (action.equals("insert")) {
    workString += doSQL(con, insertScript );
} else if (action.equals("update")) {
    workString += updater(con);
} else if (action.equals("select")) {
    workString += selecter(con);
} else if (action.equals("call")) {
    workString += caller(con);
} else { return "Invalid action: " + action ;
}
return workString;
}

```

The doSQL() method

```
// A JDBCExamples method to execute an SQL statement.

public static String doSQL (Connection con, String action)
    throws Exception {

    Statement stmt = con.createStatement();
    int res = stmt.executeUpdate(action);
    return "performed";
}
```

The updater() method

```
// A method that updates a certain row of the 'xmp' table.
// This method illustrates prepared statements and parameter markers.

public static String updater(Connection con)
    throws Exception {

    String sql = "update xmp set name = ?, home = ? where id = ?";
    int id=1;
    Address home = new Address("123 Main", "98765");
    String name = "Sam Brown";
    PreparedStatement pstmt = con.prepareStatement(sql);
    pstmt.setString(1, name);
    pstmt.setObject(2, home);
    pstmt.setInt(3, id);
    int res = pstmt.executeUpdate();
    return "performed";
}
```

The selector() method

```
// A JDBCExamples method to retrieve a certain row
// of the 'xmp' table.
// This method illustrates prepared statements, parameter markers,
// and result sets.

public static String selector(Connection con)
    throws Exception {

    String sql = "select name, home from xmp where id=?";
```

```

int id=1;
Address home = null;
String name = "";
String street = "";
String zip = "";
PreparedStatement pstmt = con.prepareStatement(sql);
pstmt.setInt(1, id);
ResultSet rs = pstmt.executeQuery();
if (rs.next()) {
    name = rs.getString(1);
    home = (Address)rs.getObject(2);
    if (rs.next()) {
        throw new Exception("Error:  Select returned
                               multiple rows");
    } else { // No action
    }
} else { throw new Exception("Error:  Select returned no rows");
}
return "- Row with id=1:  name("+ name + )
        + " street(" + home.street + ) zip("+ home.zip + );

```

The caller() method

```

// A JDBCExamples method to call a stored procedure,
// passing input and output parameters of datatype String
// and Address.
// This method illustrates callable statements, parameter markers,
// and result sets.

```

```

public static String caller(Connection con)
    throws Exception {
    CallableStatement cs = con.prepareCall("{call inoutproc
        (?, ?, ?, ?, ?)}");
    int id = 1;
    String newName = "Frank Farr";
    Address newHome = new Address("123 Farr Lane", "87654");
    cs.setInt(1, id);
    cs.setString(2, newName);
    cs.setObject(3, newHome);
    cs.registerOutParameter(4, java.sql.Types.VARCHAR);
    cs.registerOutParameter(5, java.sql.Types.JAVA_OBJECT);
    int res = cs.executeUpdate();
    String oldName = cs.getString(4);
    Address oldHome = (Address)cs.getObject(5);
    return "- Old values of row with id=1: name("+oldName+ )

```

```
        street(" + oldHome.street + ")    zip("+ oldHome.zip + );  
    }  
}
```

SQLJ Functions and Stored Procedures

This chapter describes how to wrap Java methods in SQL names and use them as Adaptive Server functions and stored procedures.

Name	Page
Overview	75
Invoking Java methods in Adaptive Server	78
Using Sybase Central to manage SQLJ functions and procedures	80
SQLJ user-defined functions	81
SQLJ stored procedures	86
Viewing information about SQLJ functions and procedures	97
Advanced topics	97
SQLJ and Sybase implementation: a comparison	102
SQLJExamples class	105

Overview

You can enclose Java static methods in SQL wrappers and use them exactly as you would Transact-SQL stored procedures or built-in functions. This functionality:

- Allows Java methods to return output parameters and result sets to the calling environment.
- Complies with Part 1 of the ANSI SQLJ standard specification.
- Allows you to take advantage of traditional SQL syntax, metadata, and permission capabilities.
- Allows you to use existing Java methods as SQLJ procedures and functions on the server, on the client, and on any SQLJ-compliant, third-party database.

❖ **Creating a SQLJ stored procedure or function**

Perform these steps to create and execute a SQLJ stored procedure or function.

- 1 Create and compile the Java method. Install the method class in the database using the `installjava` utility.

Refer to Chapter 2, “Preparing for and Maintaining Java in the Database,” for information on creating, compiling, and installing Java methods in Adaptive Server.
- 2 Using the SQLJ `create procedure` or `create function` statement, define a SQL name for the method.
- 3 Execute the procedure or function. The examples in this chapter use JDBC method calls or `isql`. You can also execute the method using Embedded SQL or ODBC.

Compliance with SQLJ Part 1 specifications

Adaptive Server SQLJ stored procedures and functions comply with SQLJ Part 1 of the standard specifications for using Java with SQL. See “Standards” on page 4 for a description of the SQLJ standards.

Adaptive Server supports most features described in the SQLJ Part 1 specification; however, there are some differences. Unsupported features are listed in Table 5-3 on page 103; partially supported features are listed in Table 5-4 on page 103. Sybase-defined features—those not defined by the standard but left to the implementation—are listed in Table 5-5 on page 103.

In those instances where Sybase proprietary implementation differs from the SQLJ specifications, Sybase supports the SQLJ standard. For example, non-Java Sybase SQL stored procedures support two parameter modes: `in` and `inout`. The SQLJ standard supports three parameter modes: `in`, `out`, and `inout`. The Sybase syntax for creating SQLJ stored procedures supports all three parameter modes.

General issues

This section describes general issues and constraints that apply to SQLJ functions and stored procedures.

Security and permissions

Sybase provides different security models for SQLJ stored procedures and SQLJ functions.

SQLJ functions and user-defined functions (UDFs) (see “Invoking Java methods in SQL” on page 28) use the same security model. Permission to execute any UDF or SQLJ function is granted implicitly to `public`. If the function performs SQL queries via JDBC, permission to access the data is checked against the invoker of the function. Thus, if user A invokes a function that accesses table `t1`, user A must have `select` permission on `t1` or the query fails.

SQLJ stored procedures use the same security model as Transact-SQL stored procedures. The user must be granted explicit permission to execute a SQLJ or Transact-SQL stored procedure. If a SQLJ procedure performs SQL queries via JDBC, implicit permission grant support is applied. This security model allows the owner of the stored procedure, if the owner owns all SQL objects referenced by the procedure, to grant `execute` permission on the procedure to another user. The user who has `execute` permission can execute all SQL queries in the stored procedure, even if the user does not have permission to access those objects.

For a more detailed description of security for stored procedures, see the *System Administration Guide*.

SQLJ Examples

The examples used in this chapter assume a SQL table called `sales_emps` with these columns:

- `name` – the employee’s name
- `id` – the employee’s identification number
- `state` – the state in which the employee is located
- `sales` – amount of the employee’s sales
- `jobcode` – the employee’s job code

The table definition is:

```
create table sales_emps
  (name varchar(50), id char(5),
   state char(20), sales decimal (6,2),
   jobcode integer null)
```

The example class is `SQLJExamples`, and the methods are:

- `region()` – maps a U.S. state code to a region number. The method does not use SQL.
- `correctStates()` – performs a SQL update command to correct the spelling of state codes. Old and new spellings are specified by input parameters.
- `bestTwoEmps()` – determines the top two employees by their sales records and returns those values as output parameters.
- `SQLJExamplesorderedEmps()` – creates a SQL result set consisting of selected employee rows ordered by values in the `sales` column, and returns the result set to the client.
- `job()` – returns a string value corresponding to an integer job code value.

See “`SQLJExamples` class” on page 105 for the text of each method.

Invoking Java methods in Adaptive Server

You can invoke Java methods in two different ways in Adaptive Server:

- Invoke Java methods directly in SQL. Directions for invoking methods in this way are presented in Chapter 3, “Using Java Classes in SQL.”
- Invoke Java methods indirectly using SQLJ stored procedures and functions that provide Transact-SQL aliases for the method name. This chapter describes invoking Java methods in this way.

Whichever way you choose, you must first create your Java methods and install them in the Adaptive Server database using the `installjava` utility. See Chapter 2, “Preparing for and Maintaining Java in the Database,” for more information.

Invoking Java methods directly with their Java names

You can invoke Java methods in SQL by referencing them with their fully qualified Java names. Reference instances for instance methods, and either instances or classes for static methods.

You can use static methods as user-defined functions (UDFs) that return a value to the calling environment. You can use a Java static method as a UDF in stored procedures, triggers, `where` clauses, `select` statements, or anywhere that you can use a built-in SQL function.

When you call a Java method using its name, you cannot use methods that return output parameters or result sets to the calling environment. A method can manipulate the data it receives from a JDBC connection, but the method can only return the single return value declared in its definition to the calling environment.

You cannot use cross-database invocations of UDF functions.

See Chapter 3, “Using Java Classes in SQL,” for information about using Java methods in this way.

Invoking Java
methods indirectly
using SQLJ

You can invoke Java methods as SQLJ functions or stored procedures. By wrapping the Java method in a SQL wrapper, you take advantage of these capabilities:

- You can use SQLJ stored procedures to return result sets and output parameters to the calling environment.
- You can take advantage of SQL metadata capabilities. For example, you can view a list of all stored procedures or functions in the database.
- SQLJ provides a SQL name for a method, which allows you to protect the method invocation with standard SQL permissions.
- Sybase SQLJ conforms to the recognized SQLJ Part 1 standard, which allows you to use Sybase SQLJ procedures and functions in conforming non-Sybase environments.
- You can invoke SQLJ functions and SQLJ stored procedures across databases.
- Because Adaptive Server checks datatype mapping when the SQLJ routine is created, you need not be concerned with datatype mapping when executing the routines.

You must reference static methods in a SQLJ routine; you cannot reference instance methods.

This chapter describes how you can use Java methods as SQLJ stored procedures and functions.

Using Sybase Central to manage SQLJ functions and procedures

You can manage SQLJ functions and procedures from the command line using `isql` and from the Adaptive Server plug-in to Sybase Central. From the Adaptive Server plug-in you can:

- Create a SQLJ function or procedure
- Execute a SQLJ function or procedure
- View and modify the properties of a SQLJ function or procedure
- Delete a SQLJ function or procedure
- View the dependencies of a SQLJ function or procedure
- Create permissions for a SQLJ procedure

The following procedures describes how to create and view the properties of a SQLJ routine. You can view dependencies and create and view permissions from the routine's property sheet.

❖ **Creating a SQLJ function/procedure**

First, create and compile the Java method. Install the method class in the database using `installjava`. Then follow these steps:

- 1 Start the Adaptive Server plug-in and connect to Adaptive Server.
- 2 Double-click on the database in which you want to create the routine.
- 3 Open the SQLJ Procedures/SQLJ Functions folder.
- 4 Double-click the Add new Java Stored Procedure/Function icon.
- 5 Use the Add new Java Stored Procedure/Function wizard to create the SQLJ procedure or function.

When you have finished using the wizard, the Adaptive Server plug-in displays the SQLJ routine you have created in an edit screen, where you can modify the routine and execute it.

- v **To view the properties of a SQLJ function or procedure**
 - 1 Start the Adaptive Server plug-in and connect to Adaptive Server.
 - 2 Double-click on the database in which the routine is stored.
 - 3 Open the SQLJ Procedures/SQLJ Functions folder.
 - 4 Highlight a function or procedure icon.
 - 5 Select File | Properties.

SQLJ user-defined functions

The `create function` command specifies a SQLJ function name and signature for a Java method. You can use SQLJ functions to read and modify SQL and to return a value described by the referenced method.

The SQLJ syntax for `create function` is:

```
create function [owner].sql_function_name
    ([sql_parameter_name sql_datatype
      [( length)| (precision[, scale])]
    [, sql_parameter_name sql_datatype
      [( length) | ( precision[, scale]) ] ]
    ...])
returns sql_datatype
    [( length)| (precision[, scale])]
[modifies sql data]
[returns null on null input |
  called on null input]
[deterministic | not deterministic]
[exportable]
language java
parameter style java
external name 'java_method_name
    [[java_datatype[ {, java_datatype }
    ...]]]'
```

When creating a SQLJ function:

- The **SQL function signature** is the SQL datatype `sql_datatype` of each function parameter.
- To comply with the ANSI standard, do not include an @ sign before parameter names.

Sybase adds an @ sign internally to support parameter name binding. You will see the @ sign when using `sp_help` to print out information about the SQLJ stored procedure.

- When creating a SQLJ function, you must include the parentheses that surround the *sql_parameter_name* and *sql_datatype* information—even if you do not include that information.

For example:

```
create function sqlj_fc()  
  language java  
  parameter style java  
  external name 'SQLJExamples.method'
```

- The `modifies sql data` clause specifies that the method invokes SQL operations and reads and modifies SQL data. This is the default value. You do not need to include it except for syntactic compatibility with the SQLJ Part 1 standard.
- `returns null on null input` and `called on null input` specify how Adaptive Server handles null arguments of a function call. `returns null on null input` specifies that if the value of any argument is null at runtime, the return value of the function is set to null and the function body is not invoked. `called on null input` is the default. It specifies that the function is invoked regardless of null argument values.

Function calls and null argument values are described in detail in “Handling nulls in the function call” on page 85.

- You can include the `deterministic` or `not deterministic` keywords, but Adaptive Server does not use them. They are included for syntactic compatibility with the SQLJ Part 1 standard.
- The `exportable` keyword specifies that the function is to run on a remote server using Sybase OmniConnect™ capabilities. Both the function and the method on which it is based must be installed on the remote server.
- The `language java` and `parameter style java` specify that the referenced method is written in Java and that the parameters are Java parameters. You *must* include these phrases when creating a SQLJ function.
- The `external name` clause specifies that the routine is not written in SQL and identifies the Java method, class and, package name (if any).

- The Java method signature specifies the Java datatype *java_datatype* of each method parameter. The Java method signature is optional. If it is not specified, Adaptive Server infers the Java method signature from the SQL function signature.

Sybase recommends that you include the method signature as this practice handles all datatype translations. See “Mapping Java and SQL datatypes” on page 97.

- You can define different SQL names for the same Java method using `create function` and then use them in the same way.

Writing the Java method

Before you can create a SQLJ function, you must write the Java method that it references, compile the method class, and install it in the database.

In this example, `SQLJExamples.region()` maps a state code to a region number and returns that number to the user.

```
public static int region(String s)
    throws SQLException {
    s = s.trim();
    if (s.equals("MN") || s.equals("VT") ||
        s.equals("NH") ) return 1;
    if (s.equals("FL") || s.equals("GA") ||
        s.equals("AL") ) return 2;
    if (s.equals("CA") || s.equals("AZ") ||
        s.equals("NV") ) return 3;
    else throw new SQLException
        ("Invalid state code", "X2001");
}
```

Creating the SQLJ function

After writing and installing the method, you can create the SQLJ function. For example:

```
create function region_of(state char(20))
    returns integer
language java parameter style java
external name
    'SQLJExamples.region(java.lang.String)'
```

The SQLJ `create function` statement specifies an input parameter (`state char(20)`) and an integer return value. The SQL function signature is `char(20)`. The Java method signature is `java.lang.String`.

Calling the function

You can call a SQLJ function directly, as if it were a built-in function. For example:

```
select name, region_of(state) as region
       from sales_emps
where region_of(state)=3
```

Note The search sequence for functions in Adaptive Server is:

- 1 Built-in functions
 - 2 SQLJ functions
 - 3 Java-SQL functions that are called directly
-

Handling null argument values

Java class datatypes and Java primitive datatypes handle null argument values in different ways.

- **Java object datatypes** that are classes—such as `java.lang.Integer`, `java.lang.String`, `java.lang.byte[]`, and `java.sql.Timestamp`—can hold both actual values and null reference values.
- **Java primitive datatypes**—such as `boolean`, `byte`, `short`, and `int`—have no representation for a null value. They can hold only non-null values.

When a Java method is invoked that causes a SQL null value to be passed as an argument to a Java parameter whose datatype is a Java class, it is passed as a Java null reference value. When a SQL null value is passed as an argument to a Java parameter of a Java primitive datatype, however, an exception is raised because the Java primitive datatype has no representation for a null value.

Typically, you will write Java methods that specify Java parameter datatypes that are classes. In this case, nulls are handled without raising an exception. If you choose to write Java functions that use Java parameters that cannot handle null values, you can either:

- Include the `returns null on null input` clause when you create the SQLJ function, or
- Invoke the SQLJ function using a `case` or other conditional expression to test for null values and call the SQLJ function only for the non-null values.

You can handle expected nulls when you create the SQLJ function or when you call it. The following sections describe both scenarios, and reference this method:

```
public static String job(int jc)
```

```
        throws SQLException {  
    if (jc==1) return "Admin";  
    else if (jc==2) return "Sales";  
    else if (jc==3) return "Clerk";  
    else return "unknown jobcode";  
}
```

Handling nulls when creating the function

If null values are expected, you can include the `returns null on null input` clause when you create the function. For example:

```
create function job_of(jc integer)  
    returns varchar(20)  
returns null on null input  
language java parameter style java  
external name 'SQLJExamples.job(int)'
```

You can then call `job_of` in this way:

```
select name, job_of(jobcode)  
    from sales_emp  
where job_of(jobcode) <> "Admin"
```

When the SQL system evaluates the call `job_of(jobcode)` for a row of `sales_emps` in which the `jobcode` column is null, the value of the call is set to null without actually calling the Java method `SQLJExamples.job`. For rows with non-null values of the `jobcode` column, the call is performed normally.

Thus, when a SQLJ function created using the `returns null on null input` clause encounters a null argument, the result of the function call is set to null and the function is not invoked.

Note If you include the `returns null on null input` clause when creating a SQLJ function, the `returns null on null input` clause applies to *all* function parameters, including nullable parameters.

If you include the `called on null input` clause (the default), null arguments for non-nullable parameters generates an exception.

Handling nulls in the function call

You can use a conditional function call to handle null values for non-nullable parameters. The following example uses a `case` expression:

```
select name,  
       case when jobcode is not null  
         then job_of(jobcode)  
         else null end  
from sales_ems where  
       case when jobcode is not null  
         then job_of(jobcode)  
         else null end <> "Admin"
```

In this example, we assume that the function `job_of` was created using the default clause called on null input.

Deleting a SQLJ function name

You can delete the SQLJ function name for a Java method using the `drop function` command. For example, enter:

```
drop function region_of
```

which deletes the `region_of` function name and its reference to the `SQLJExamples.region` method. `drop function` does not affect the referenced Java method or class.

See the *Reference Manual* for complete syntax and usage information.

SQLJ stored procedures

Using Java-SQL capabilities, you can install Java classes in the database and then invoke those methods from a client or from within the SQL system. You can also invoke Java static (class) methods in another way—as SQLJ stored procedures.

SQLJ stored procedures:

- Can return result sets and/or output parameters to the client
- Behave exactly as Transact-SQL stored procedures when executed
- Can be called from the client using ODBC, isql, or JDBC
- Can be called within the server from other stored procedures or native Adaptive Server JDBC

The end user need not know whether the procedure being called is a SQLJ stored procedure or a Transact-SQL stored procedure. They are both invoked in the same way.

The SQLJ syntax for create procedure is:

```
create procedure [owner.]sql_procedure_name
  ([[ in | out | inout ] sql_parameter_name
    sql_datatype [( length) |
    (precision[, scale])])
  [, [ in | out | inout ] sql_parameter_name
    sql_datatype [( length) |
    (precision[, scale]) ]]
  ...])
[modifies sql data]
[dynamic result sets integer]
[deterministic | not deterministic]
language java
parameter style java
external name 'java_method_name
  ([[java_datatype[, java_datatype
  ...]])]'
```

Note To comply with the ANSI standard, the SQLJ create procedure command syntax is different from syntax used to create Sybase Transact-SQL stored procedures.

Refer to the *Reference Manual* for a detailed description of each keyword and option in this command.

When creating SQLJ stored procedures:

- The **SQL procedure signature** is the SQL datatype *sql_datatype* of each procedure parameter.
- When creating a SQLJ stored procedure, do not include an @ sign before parameter names. This practise is compliant with the ANSI standard.

Sybase adds an @ sign internally to support parameter name binding. You will see the @ sign when using `sp_help` to print out information about the SQLJ stored procedure.

- When creating a SQLJ stored procedure, you must include the parentheses that surround the *sql_parameter_name* and *sql_datatype* information—even if you do not include that information.

For example:

```
create procedure sqlj_sproc ()  
    language java  
    parameter style java  
    external name "SQLJExamples.method1"
```

- You can include the keywords `modifies sql data` to indicate that the method invokes SQL operations and reads and modifies SQL data. This is the default value.
- You must include the dynamic result sets *integer* option when result sets are to be returned to the calling environment. Use the *integer* variable to specify the maximum number of result sets expected.
- You can include the keywords `deterministic` or `not deterministic` for compatibility with the SQLJ standard. However, Adaptive Server does not make use of this option.
- You must include the `language java` parameter and `style java` keywords, which tell Adaptive Server that the external routine is written in Java and the runtime conventions for arguments passed to the external routine are Java conventions.
- The `external name` clause indicates that the external routine is written in Java and identifies the Java method, class, and package name (if any).
- The Java method signature specifies the Java datatype *java_datatype* of each method parameter. The Java method signature is optional. If one is not specified, Adaptive Server infers one from the SQL procedure signature.

Sybase recommends that you include the method signature as this practice handles all datatype translations. See “Mapping Java and SQL datatypes” on page 97 for more information.

- You can define different SQL names for the same Java method using `create procedure` and then use them in the same way.

Modifying SQL data

You can use a SQLJ stored procedure to modify information in the database. The method referenced by the SQLJ procedure must be either:

- A method of type `void`, or
- A method with an `int` return type (incorporation of the `int` return type is a Sybase extension of the SQLJ standard).

Writing the Java method

The method `SQLJExamples.correctStates()` performs a SQL update statement to correct the spelling of state codes. Input parameters specify the old and new spellings. `correctStates()` is a void method; no value is returned to the caller.

```
public static void correctStates(String oldSpelling,
                                String newSpelling) throws SQLException {

    Connection conn = null;
    PreparedStatement pstmt = null;
    try {
        Class.forName("sybase.asejdbc.ASEDriver");
        conn = DriverManager.getConnection
            ("jdbc:default:connection");
    }
    catch (Exception e) {
        System.err.println(e.getMessage() +
            ":error in connection");
    }
    try {
        pstmt = conn.prepareStatement
            ("UPDATE sales_emps SET state = ?
             WHERE state = ?");
        pstmt.setString(1, newSpelling);
        pstmt.setString(2, oldSpelling);
        pstmt.executeUpdate();
    }
    catch (SQLException e) {
        System.err.println("SQLException: " +
            e.getErrorCode() + e.getMessage());
    }
    return;
}
```

Creating the stored procedure

Before you can call a Java method with a SQL name, you must create the SQL name for it using the SQLJ create procedure command. The `modifies sql data` clause is optional.

```
create procedure correct_states(old char(20),
                                not_old char(20))
modifies sql data
language java parameter style java
external name
    'SQLJExamples.correctStates
    (java.lang.String, java.lang.String)'
```

The `correct_states` procedure has a SQL procedure signature of `char(20), char(20)`. The Java method signature is `java.lang.String, java.lang.String`.

Calling the stored procedure

You can execute the SQLJ procedure exactly as you would a Transact-SQL procedure. In this example, the procedure executes from isql:

```
execute correct_states 'GEO', 'GA'
```

Using input and output parameters

Java methods do not support output parameters. When you wrap a Java method in SQL, however, you can take advantage of Sybase SQLJ capabilities that allow input, output, and input/output parameters for SQLJ stored procedures.

When you create a SQLJ procedure, you identify the mode for each parameter as `in`, `out`, or `inout`.

- For input parameters, use the `in` keyword to qualify the parameter. `in` is the default; Adaptive Server assumes an input parameter if you do not enter a parameter mode.
- For output parameters, use the `out` keyword.
- For parameters that can pass values both to and from the referenced Java method, use the `inout` keyword.

Note You create Transact-SQL stored procedures using only the `in` and `out` keywords. The `out` keyword corresponds to the SQLJ `inout` keyword. See the create procedure reference pages in the *Adaptive Server Reference Manual* for more information.

To create a SQLJ stored procedure that defines output parameters, you must:

- Define the output parameter(s) using either the `out` or `inout` option when you create the SQLJ stored procedure.
- Declare those parameters as Java arrays in the Java method. SQLJ uses arrays as containers for the method's output parameter values.

For example, if you want an Integer parameter to return a value to the caller, you must specify the parameter type as `Integer[]` (an array of Integer) in the method.

The array object for an out or inout parameter is created implicitly by the system. It has a single element. The input value (if any) is placed in the first (and only) element of the array before the Java method is called. When the Java method returns, the first element is removed and assigned to the output variable. Typically, this element will be assigned a new value by the called method.

The following examples illustrate the use of output parameters using a Java method `bestTwoEmps()` and a stored procedure `best2` that references that method.

Writing the Java method

The `SQLJExamples.bestTwoEmps()` method returns the name, ID, region, and sales of the two employees with the highest sales performance records. The first eight parameters are output parameters requiring a containing array. The ninth parameter is an input parameter and does not require an array.

```
public static void bestTwoEmps(String[] n1,
    String[] id1, int[] r1,
    BigDecimal[] s1, String[] n2,
    String[] id2, int[] r2, BigDecimal[] s2,
    int regionParm) throws SQLException {

    n1[0] = "*****";
    id1[0] = "";
    r1[0] = 0;
    s1[0] = new BigDecimal(0);
    n2[0] = "*****",
    id2[0] = "";
    r2[0] = 0;
    s2[0] = new BigDecimal(0);

    try {
        Connection conn = DriverManager.getConnection
            ("jdbc:default:connection");
        java.sql.PreparedStatement stmt =
            conn.prepareStatement("SELECT name, id,"
                + "region_of(state) as region, sales FROM"
                + "sales_emps WHERE"
                + "region_of(state)>? AND"
                + "sales IS NOT NULL ORDER BY sales DESC");
        stmt.setInt(1, regionParm);
        ResultSet r = stmt.executeQuery();

        if(r.next()) {
            n1[0] = r.getString("name");
            id1[0] = r.getString("id");
            r1[0] = r.getInt("region");
```

```
        s1[0] = r.getBigDecimal("sales");
    }
    else return;

    if(r.next()) {
        n2[0] = r.getString("name");
        id2[0] = r.getString("id");
        r2[0] = r.getInt("region");
        s2[0] = r.getBigDecimal("sales");
    }
    else return;
}
catch (SQLException e) {
    System.err.println("SQLException: " +
        e.getErrorCode() + e.getMessage());
}
}
```

Creating the SQLJ procedure

Create a SQL name for the `bestTwoEmps` method. The first eight parameters are output parameters; the ninth is an input parameter.

```
create procedure best2
(out n1 varchar(50), out id1 varchar(5),
 out s1 decimal(6,2), out r1 integer,
 out n2 varchar(50), out id2 varchar(50),
 out r2 integer, out s2 decimal(6,2),
 in region integer)
language java
parameter style java
external name
'SQLJExamples.bestTwoEmps (java.lang.String,
 java.lang.String, int, java.math.BigDecimal,
 java.lang.String, java.lang.String, int,
 java.math.BigDecimal, int)'
```

The SQL procedure signature for `best2` is: `varchar(20), varchar(5), decimal(6,2)` and so on. The Java method signature is `String, String, int, BigDecimal` and so on.

Calling the procedure

After the method is installed in the database and the SQLJ procedure referencing the method has been created, you can call the SQLJ procedure.

At runtime, the SQL system:

- 1 Creates the needed arrays for the `out` and `inout` parameters when the SQLJ procedure is called.

- 2 Copies the contents of the parameter arrays into the out and inout target variables when returning from the SQLJ procedure.

The following example calls the `best2` procedure from `isql`. The value for the `region` input parameter specifies the region number.

```
declare @n1 varchar(50), @id1 varchar(5),
        @s1 decimal (6,2), @r1 integer, @n2 varchar(50),
        @id2 varchar(50), @r2 integer, @s2 decimal(6,2),
        @region integer
select @region = 3
execute best2 @n1 out, @id1 out, @s1 out, @r1 out,
            @n2 out, @id2 out, @r2 out, @s2 out, @region
```

Note Adaptive Server calls SQLJ stored procedures exactly as it calls Transact-SQL stored procedures. Thus, when using `isql` or any other non-Java client, you must precede parameter names by the `@` sign.

Returning result sets

A SQL result set is a sequence of SQL rows that is delivered to the calling environment.

When a Transact-SQL stored procedure returns one or more results sets, those result sets are implicit output from the procedure call. That is, they are not declared as explicit parameters or return values.

Java methods can return Java result set objects, but they do so as explicitly declared method values.

To return a SQL-style result set from a Java method, you must first wrap the Java method in a SQLJ stored procedure. When you call the method as a SQLJ stored procedure, the result sets, which are returned by the Java method as Java result set objects, are transformed by the server to SQL result sets.

When writing the Java method to be invoked as a SQLJ procedure that returns a SQL-style result set, you must specify an additional parameter to the method for each result set that the method can return. Each such parameter is a single-element array of the Java `ResultSet` class.

This section describes the basic process of writing a method, creating the SQLJ stored procedure, and calling the method. See “Specifying Java method signatures explicitly or implicitly” on page 99 for more information about returning result sets.

Writing the Java method

The following method, `SQLJExamples.orderedEmps`, invokes SQL, includes a `ResultSet` parameter, and uses JDBC calls for securing a connection and opening a statement.

```
public static void orderedEmps
    (int regionParm, ResultSet[] rs) throws
        SQLException {

    Connection conn = null;
    PreparedStatement pstmt = null;

    try {
        Class.forName
            ("sybase.asejdbc.ASEDriver");
        Connection conn =
            DriverManager.getConnection
                ("jdbc:default:connection");
    }
    catch (Exception e) {
        System.err.println(e.getMessage()
            + ":error in connection");
    }

    try {
        java.sql.PreparedStatement
            stmt = conn.prepareStatement
                ("SELECT name, region_of(state) "
                "as region, sales FROM sales_emp"
                "WHERE region_of(state) > ? AND "
                "sales IS NOT NULL"
                "ORDER BY sales DESC");
        stmt.setInt(1, regionParm);
        rs[0] = stmt.executeQuery();
        return;
    }
    catch (SQLException e)
        System.err.println("SQLException: "
            + e.getErrorCode() + e.getMessage());
    }
    return;
}
```

`orderedEmps` returns a single result set. You can also write methods that return multiple result sets. For each result set returned, you must:

- Include a separate `ResultSet` array parameter in the method signature.
- Create a `Statement` object for each result set.
- Assign each result set to the first element of its `ResultSet` array.

Adaptive Server always returns the current open `ResultSet` object for each `Statement` object. When creating Java methods that return result sets:

- Create a `Statement` object for each result set that is to be returned to the client.
- Do not explicitly close `ResultSet` and `Statement` objects. Adaptive Server closes them automatically.

Note Adaptive Server ensures that `ResultSet` and `Statement` objects are not closed by garbage collection unless and until the affected result sets have been processed and returned to the client.

- If some rows of the result set are fetched by calls of the Java `next()` method, only the remaining rows of the result set are returned to the client.

Creating the SQLJ stored procedure

When you create a SQLJ stored procedure that returns result sets, you must specify the maximum number of result sets that can be returned. In this example, the `ranked_emps` procedure returns a single result set.

```
create procedure ranked_emps(region integer)
dynamic result sets 1
language java parameter style java
external name 'SQLJExamples.orderedEmps(int,
    ResultSet [])'
```

If `ranked_emps` generates more result sets than are specified by `create procedure`, a warning displays and the procedure returns only the number of result sets specified. As written, the `ranked_emps` SQLJ stored procedures matches only one Java method.

Note Some restrictions apply to method overloading when you infer a method signature involving result sets. See “Mapping Java and SQL datatypes” on page 97 for more information.

Calling the procedure

After you have installed the method’s class in the database and created the SQLJ stored procedure that references the method, you can call the procedure. You can write the call using any mechanism that processes SQL result sets.

For example, to call the `ranked_emps` procedure using JDBC, enter the following:

```
java.sql.CallableStatement stmt =
    conn.prepareCall("{call ranked_emps(?)}");
stmt.setInt(1,3);
ResultSet rs = stmt.executeQuery();
while (rs.next()) {
    String name = rs.getString(1);
    int region = rs.getInt(2);
    BigDecimal sales = rs.getBigDecimal(3);
    System.out.print("Name = " + name);
    System.out.print("Region = " + region);
    System.out.print("Sales = " + sales);
    System.out.println();
}
```

The `ranked_emps` procedure supplies only the parameter declared in the `create procedure` statement. The SQL system supplies an empty array of `ResultSet` parameters and calls the Java method, which assigns the output result set to the array parameter. When the Java method completes, the SQL system returns the result set in the output array element as a SQL result set.

Note You can return result sets from a temporary table only when using an external JDBC driver such as `jConnect`. You cannot use the Adaptive Server native JDBC driver for this task.

Deleting a SQLJ stored procedure name

You can delete the SQLJ stored procedure name for a Java method using the `drop procedure` command. For example, enter:

```
drop procedure correct_states
```

which deletes the `correct_states` procedure name and its reference to the `SQLJExamples.correctStates` method. `drop procedure` does not affect the Java class and method referenced by the procedure.

Viewing information about SQLJ functions and procedures

Several system stored procedures can provide information about SQLJ routines:

- `sp_depends` lists database objects referenced by the SQLJ routine and database objects that reference the SQLJ routine.
- `sp_help` lists each parameter name, type, length, precision, scale, parameter order, parameter mode and return type of the SQLJ routine.
- `sp_helpjava` lists information about Java classes and JARs installed in the database. The `depends` parameter lists dependencies of specified classes that are named in the `external name` clause of the SQLJ `create function` or `SQLJ create procedure` statement.
- `sp_helprotect` reports the permissions of SQLJ stored procedures and SQLJ functions.

See the *Adaptive Server Reference Manual* for complete syntax and usage information for these system procedures.

Advanced topics

The following topics present a detailed description of SQLJ topics for advanced users.

Mapping Java and SQL datatypes

When you create a stored procedure or function that references a Java method, the datatypes of input and output parameters or result sets must not conflict when values are converted from the SQL environment to the Java environment and back again. The rules for how this mapping takes place are consistent with the JDBC standard implementation. They are shown below and in Table 5-1 on page 98.

Each SQL parameter and its corresponding Java parameter must be mappable. SQL and Java datatypes are mappable in these ways:

- A SQL datatype and a primitive Java datatype are *simply mappable* if so specified in Table 5-1.
- A SQL datatype and a non-primitive Java datatype are *object mappable* if so specified in Table 5-1.
- A SQL abstract datatype (ADT) and a non-primitive Java datatype are *ADT mappable* if both are the same class or interface.
- A SQL datatype and a Java datatype are *output mappable* if the Java datatype is an array and the SQL datatype is simply mappable, object mappable, or ADT mappable to the Java datatype. For example, character and `String[]` are output mappable.
- A Java datatype is *result-set mappable* if it is an array of the result set-oriented class: `java.sql.ResultSet`.

In general, a Java method is mappable to SQL if each of its parameters is mappable to SQL and its result set parameters are result-set mappable and the return type is either mappable (functions) or void or int (procedures).

Support for int return types for SQLJ stored procedures is a Sybase extension of the SQLJ Part 1 standard.

Table 5-1: Simply and object mappable SQL and Java datatypes

SQL datatype	Corresponding Java datatypes	
	Simply mappable	Object mappable
char/unichar		java.lang.String
nchar		java.lang.String
varchar/univarchar		java.lang.String
nvarchar		java.lang.String
text		java.lang.String
numeric		java.math.BigDecimal
decimal		java.math.BigDecimal
money		java.math.BigDecimal
smallmoney		java.math.BigDecimal
bit	boolean	Boolean
tinyint	byte	Integer
smallint	short	Integer
integer	int	Integer
bigint	long	java.math.BigInteger
unsigned smallint	int	Integer
unsigned int	long	Integer

SQL datatype	Corresponding Java datatypes	
	Simply mappable	Object mappable
unsigned bigint		java.math.BigInteger
real	float	Float
float	double	Double
double precision	double	Double
binary		byte[]
varbinary		byte[]
datetime		java.sql.Timestamp
smalldatetime		java.sql.Timestamp
date		java.sql.Date
time		java.sql.Time

Specifying Java method signatures explicitly or implicitly

When you create a SQLJ function or stored procedure, you typically specify a Java method signature. You can also allow Adaptive Server to infer the Java method signature from the routine's SQL signature according to standard JDBC datatype correspondence rules described earlier in this section and in Table 5-1.

Sybase recommends that you include the Java method signature as this practise ensures that all datatype translations are handled as specified.

You can allow Adaptive Server to infer the method signature for datatypes that are:

- Simply mappable
- ADT mappable
- Output mappable
- Result-set mappable

For example, if you want Adaptive Server to infer the method signature for `correct_states`, the create procedure statement is:

```
create procedure correct_states(old char(20),
                             not_old char(20))
modifies sql data
language java parameter style java
external name 'SQLJExamples.correctStates'
```

Adaptive Server infers a Java method signature of `java.lang.String` and `java.lang.String`. If you explicitly add the Java method signature, the create procedure statement looks like this:

```
create procedure correct_states(old char(20),
                               not_old char(20))
  modifies sql data
  language java parameter style java
  external name 'SQLJExamples.correctStates'
  (java.lang.String, java.lang.String)'
```

You *must* explicitly specify the Java method signature for datatypes that are object mappable. Otherwise, Adaptive Server infers the primitive, simply mappable datatype.

For example, the `SQLJExamples.job` method contains a parameter of type `int`. (See “Handling null argument values” on page 84.) When creating a function referencing that method, Adaptive Server infers a Java signature of `int`, and you need not specify it.

However, suppose the parameter of `SQLJExamples.job` was Java `Integer`, which is the object-mappable type. For example:

```
public class SQLJExamples {
    public static String job(Integer jc)
        throws SQLException ...
}
```

Then, you must specify the Java method signature when you create a function that references it:

```
create function job_of(jc integer)
...
external name
    'SQLJExamples.job(java.lang.Integer)'
```

Returning result sets
and method
overloading

When you create a SQLJ stored procedure that returns result sets, you specify the maximum number of result sets that can be returned.

If you specify a Java method signature, Adaptive Server looks for the single method that matches the method name and signature. For example:

```
create procedure ranked_emps(region integer)
  dynamic result sets 1
  language java parameter style java
  external name 'SQLJExamples.orderedEmps'
  (int, java.sql.ResultSet[])'
```

In this case, Adaptive Server resolves parameter types using normal Java overloading conventions.

Suppose, however, that you do not specify the Java method signature:

```
create procedure ranked_emps(region integer)
  dynamic result sets 1
```

```
language java parameter style java
external name 'SQLJExamples.orderedEmps'
```

If two methods exist, one with a signature of `int, RS[]`, the other with a signature of `int, RS[], RS[]`, Application Server cannot distinguish between the two methods and the procedure fails. If you allow Adaptive Server to infer the Java method signature when returning result sets, make sure that *only one method* satisfies the inferred conditions.

Note The number of dynamic result sets specified only affects the maximum number of results that can be returned. It does not affect method overloading.

Ensuring signature
validity

If an installed class has been modified, Adaptive Server checks to make sure that the method signature is valid when you invoke a SQLJ procedure or function that references that class. If the signature of a modified method is still valid, the execution of the SQLJ routine succeeds.

Using the command main method

In a Java client, you typically begin Java applications by running the Java Virtual Machine (VM) on the command `main` method of a class. The `JDBCExamples` class, for example, contains a `main` method. It is the command `main` method that executes when you execute the class from the command line as in the following:

```
java JDBCExamples
```

Note You cannot reference a Java `main` method in a SQLJ `create function` statement.

If you reference a Java `main` method in a SQLJ `create procedure` statement, the command `main` method must have the Java method signature `String[]` as in:

```
public static void main(java.lang.String[]) {
    ...
}
```

If the Java method signature is specified in the `create procedure` statement, it must be specified as `(java.lang.String[])`. If the Java method signature is not specified, it is assumed to be `(java.lang.String[])`.

If the SQLJ procedure signature contains parameters, those parameters must be char, unichar, varchar, or univarchar. At runtime, they are passed as a Java array of java.lang.String.

Each argument you provide to the SQLJ procedure must be char, unichar, varchar, univarchar, or a literal string because it is passed to the main method as an element of the java.lang.String array. You cannot use the dynamic result sets clause when creating a main procedure.

SQLJ and Sybase implementation: a comparison

This section describes differences between SQLJ Part 1 standard specifications and the Sybase proprietary implementation for SQLJ stored procedures and functions.

Table 5-2 describes Adaptive Server enhancements to the SQLJ implementation.

Table 5-2: Sybase enhancements

Category	SQLJ standard	Sybase implementation
create procedure command	Supports only Java methods that do not return values. The methods must have void return type.	Supports Java methods that allow an integer value return. The methods referenced in create procedure can have either void or integer return types.
create procedure and create function commands	Supports only SQL datatypes in create procedure or create function parameter list.	Supports SQL datatypes and nonprimitive Java datatypes as abstract data types (ADTs).
SQLJ function and SQLJ procedure invocation	Does not support implicit SQL conversion to SQLJ datatypes.	Supports implicit SQL conversion to SQLJ datatypes.
SQLJ functions	Does not allow SQLJ functions to run on remote servers.	Allows SQLJ functions to run on remote servers using Sybase OmniConnect capabilities.
drop procedure and drop function commands	Requires complete command name: drop procedure or drop function.	Supports complete function name and abridged names: drop proc and drop func.

Table 5-3 describes SQLJ standard features not included in the Sybase implementation.

Table 5-3: SQLJ features not supported

SQLJ category	SQLJ standard	Sybase implementation
create function command	Allows users to specify the same SQL name for multiple SQLJ functions.	Requires unique names for all stored procedure and functions.
utilities	Supports <code>sqlj.install_jar</code> , <code>sqlj.replace_jar</code> , <code>sqlj.remove_jar</code> , and similar utilities to install, replace, and remove JAR files.	Supports the <code>installjava</code> utility and the <code>remove java Transact-SQL</code> command to perform similar functions.

Table 5-4 describes the SQLJ standard features supported in part by the Sybase implementation.

Table 5-4: SQLJ features partially supported

SQLJ category	SQLJ standard	Sybase implementation
create procedure and create function commands	Allows users to install different classes with the same name in the same database if they are in different JAR files.	Requires unique class names in the same database.
create procedure and create function commands	Supports the key words <code>no sql</code> , contains <code>sql</code> , reads <code>sql</code> data, and modifies <code>sql</code> data to specify the SQL operations the Java method can perform.	Supports modifies <code>sql</code> data only.
create procedure command	Supports <code>java.sql.ResultSet</code> and the SQL/OLB iterator declaration.	Supports <code>java.sql.ResultSet</code> only.
drop procedure and drop function commands	Supports the key word <code>restrict</code> , which requires the user to drop all SQL objects (tables, views, and routines) that invoke the procedure or function before dropping the procedure or function.	Does not support the <code>restrict</code> key word and functionality.

Table 5-5 describes the SQLJ implementation-defined features in the Sybase implementation.

Table 5-5: SQLJ features defined by the implementation

SQLJ category	SQLJ standard	Sybase implementation
create procedure and create function commands	Supports the <code>deterministic</code> <code>not deterministic</code> keywords, which specify whether or not the procedure or function always returns the same values for the <code>out</code> and <code>inout</code> parameters and the function result.	Supports only the syntax for <code>deterministic</code> <code>not deterministic</code> , not the functionality.

SQLJ category	SQLJ standard	Sybase implementation
create procedure and create function commands	The validation of the mapping between the SQL signature and the Java method signature can be performed either when the <code>create</code> command is executed or when the procedure or function is invoked. The implementation defines when the validation is performed.	If the referenced class has been changed, performs all validations when the <code>create</code> command is executed, which enables faster execution.
create procedure and create function commands	Can specify the <code>create procedure</code> or <code>create function</code> commands within deployment descriptor files or as SQL DDL statements. The implementation defines which way (or ways) the commands are supported.	Supports <code>create procedure</code> and <code>create function</code> as SQL DDL statements outside of deployment descriptors.
Invoking SQLJ routines	When a Java method executes a SQL statement, any exception conditions are raised in the Java method as a Java exception of the <code>Exception.sql/Exception</code> subclass. The effect of the exception condition is defined by the implementation.	Follows the rules for Adaptive Server JDBC.
Invoking SQLJ routines	The implementation defines whether a Java method called using a SQL name executes with the privileges of the user who created the procedure or function or those of the invoker of the procedure or function.	SQLJ procedures and functions inherit the security features of SQL stored procedures and Java-SQL functions, respectively.
drop procedure and drop function commands	Can specify the <code>drop procedure</code> or <code>drop function</code> commands within deployment descriptor files or as SQL DDL statements. The implementation defines which way (or ways) the commands are supported.	Supports <code>create procedure</code> and <code>create function</code> as SQL DDL statements outside of deployment descriptors.

SQLJExamples class

This section displays the SQLJExamples class used to illustrate SQLJ stored procedures and functions. They are also in

\$SYBASE/\$SYBASE_ASE/sample/JavaXML/JavaXml.zip (UNIX) or
%SYBASE%\Ase-15_0\sample\JavaXML\JavaXml.zip (Windows NT).

```
import java.lang.*;
import java.sql.*;
import java.math.*;

static String _url = "jdbc:default:connection";

public class SQLExamples {

    public static int region(String s)
        throws SQLException {
        s = s.trim();
        if (s.equals("MN") || s.equals("VT") ||
            s.equals("NH") ) return 1;
        if (s.equals("FL") || s.equals("GA") ||
            s.equals("AL") ) return 2;
        if (s.equals("CA") || s.equals("AZ") ||
            s.equals("NV") ) return 3;
        else throw new SQLException
            ("Invalid state code", "X2001");
    }

    public static void correctStates
        (String oldSpelling, String newSpelling)
        throws SQLException {

        Connection conn = null;
        PreparedStatement pstmt = null;
        try {
            Class.forName
                ("sybase.asejdbc.ASEDriver");
            conn = DriverManager.getConnection(_url);
        }
        catch (Exception e) {
            System.err.println(e.getMessage() +
                ":error in connection");
        }
        try {
            pstmt = conn.prepareStatement
                ("UPDATE sales_ems SET state = ?
                 WHERE state = ?");
```

```
        pstmt.setString(1, newSpelling);
        pstmt.setString(2, oldSpelling);
        pstmt.executeUpdate();
    }
    catch (SQLException e) {
        System.err.println("SQLException: " +
            e.getErrorCode() + e.getMessage());
    }
}

public static String job(int jc)
    throws SQLException {
    if (jc==1) return "Admin";
    else if (jc==2) return "Sales";
    else if (jc==3) return "Clerk";
    else return "unknown jobcode";
}

public static String job(int jc)
    throws SQLException {
    if (jc==1) return "Admin";
    else if (jc==2) return "Sales";
    else if (jc==3) return "Clerk";
    else return "unknown jobcode";
}

public static void bestTwoEmps(String[] n1,
    String[] id1, int[] r1,
    BigDecimal[] s1, String[] n2,
    String[] id2, int[] r2, BigDecimal[] s2,
    int regionParm) throws SQLException {

    n1[0] = "*****";
    id1[0] = "";
    r1[0] = 0;
    s1[0] = new BigDecimal(0);
    n2[0] = "*****";
    id2[0] = "";
    r2[0] = 0;
    s2[0] = new BigDecimal(0);

    try {
        Connection conn = DriverManager.getConnection
            ("jdbc:default:connection");
        java.sql.PreparedStatement stmt =
            conn.prepareStatement("SELECT name, id, "
                + "region_of(state) as region, sales FROM "
                + "sales_ems WHERE"
```

```
        + "region_of(state)>? AND"
        + "sales IS NOT NULL ORDER BY sales DESC");
stmt.setInteger(1, regionParm);
ResultSet r = stmt.executeQuery();

if(r.next()) {
    n1[0] = r.getString("name");
    id1[0] = r.getString("id");
    r1[0] = r.getInt("region");
    s1[0] = r.getBigDecimal("sales");
}
else return;

if(r.next()) {
    n2[0] = r.getString("name");
    id2[0] = r.getString("id");
    r2[0] = r.getInt("region");
    s2[0] = r.getBigDecimal("sales");
}
else return;
}
catch (SQLException e) {
    System.err.println("SQLException: "+
        e.getErrorCode() + e.getMessage());
}
}

public static void orderedEmps
    (int regionParm, ResultSet[] rs) throws
    SQLException {

    Connection conn = null;
    PreparedStatement pstmt = null;

    try {
        Class.forName
            ("sybase.asejdbc.ASEDriver");
        Connection conn =
            DriverManager.getConnection
                ("jdbc:default:connection");
    }
    catch (Exception e) {
        System.err.println(e.getMessage()
            + ":error in connection");
    }
}
```

```
        try {
            java.sql.PreparedStatement
                stmt = conn.prepareStatement
                    ("SELECT name, region_of(state)"
                     "as region, sales FROM sales_emps"
                     "WHERE region_of(state) > ? AND"
                     "sales IS NOT NULL"
                     "ORDER BY sales DESC");
            stmt.setInt(1, regionParm);
            rs[0] = stmt.executeQuery();
            return;
        }
        catch (SQLException e) {
            System.err.println("SQLException:"
                               + e.getErrorCode() + e.getMessage());
        }
        return;
    }
}
```

Debugging Java in the Database

This chapter describes the Sybase Java debugger and how you can use it when developing Java in Adaptive Server.

Name	Page
Introduction to debugging Java	109
Using the debugger	110
A debugging tutorial	117

Introduction to debugging Java

You can use the Sybase Java debugger to test Java classes and fix problems with them.

How the debugger works

The Sybase Java debugger is a Java application that runs on a client machine. It connects to the database using the Sybase jConnect JDBC driver.

The debugger debugs classes running in the database. You can step through the source code for the files as long as you have the Java source code on the disk of your client machine. (Remember, the compiled classes are installed in the database, but the source code is not).

Requirements for using the Java debugger

To use the Java debugger, you need:

- A Java runtime environment such as the Sun Microsystems Java Runtime Environment, or the full Sun Microsystems JDK on your machine.

- The source code for your application on your client machine.

What you can do with the debugger

Using the Sybase Java debugger, you can:

- Trace execution – Step line by line through the code of a class running in the database. You can also look up and down the stack of functions that have been called.
- Set breakpoints – Run the code until you hit a breakpoint, and stop at that point in the code.
- Set break conditions – Breakpoints include lines of code, but you can also specify conditions when the code is to break. For example, you can stop at a line the tenth time it is executed, or only if a variable has a particular value. You can also stop whenever a particular exception is thrown in the Java application.
- Browse classes – You can browse through the classes installed into the database that the server is currently using.
- Inspect and set variables – You can inspect the values of variables alter their value when the execution is stopped at a breakpoint.
- Inspect and break on expressions – You can inspect the value of a wide variety of expressions.

Using the debugger

This section describes how to use the Java debugger. The next section provides a simple tutorial.

Starting the debugger and connecting to the database

The debugger is the JAR file *Debug.jar*, installed in your Adaptive Server installation directory in `$SYBASE/$SYBASE_ASE/debugger`. If it is not already present, add this file as the first element to your CLASSPATH environment variable.

Debug.jar contains many classes. To start the debugger you invoke the `sybase.vm.Debug` class, which has a `main()` method. You can start the debugger in three ways:

- Run the *jdebug* script located in `$SYBASE/$SYBASE_ASE/debugger`.
“A debugging tutorial” on page 117 provides a sample debugging session using the *jdebug* script.
- From the command line, enter:

```
java sybase.vm.Debug
```

In the Connect window, enter a URL, user login name, and password to connect to the database.
- From Sybase Central:
 - a Start Sybase Central and open the Utilities folder, under Adaptive Server Enterprise.
 - b Double-click the Java debugger icon in the right panel.
 - c In the Connect window, enter a URL, user login name, and password to connect to the database.

Compiling classes for debugging

Java compilers such as the Sun Microsystems `javac` compiler can compile Java classes at different levels of optimization. You can opt to compile Java code so that information used by debuggers is retained in the compiled class files.

If you compile your source code without using switches for debugging, you can still step through code and use breakpoints. However, you cannot inspect the values of local variables.

To compile classes for debugging using the `javac` compiler, use the `-g` option:

```
javac -g ClassName.java
```

Attaching to a Java VM

When you connect to a database from the debugger, the Connection window shows all currently active Java VMs under the user login name. If there are none, the debugger goes into *wait mode*. Wait mode works like this:

- Each time a new Java VM is started, it shows up in the list.
- You may choose either to debug the new Java VM or to wait for another one to appear.
- Once you have passed on a Java VM, you lose your chance to debug that Java VM. If you then decide to attach to the passed Java VM, you must disconnect from the database and reconnect. At this time, the Java VM appears as active, and you can attach to it.

The Source window

The Source window:

- Displays Java source code, with line numbers and breakpoint indicators (an asterisk in the left column).
- Displays execution status in the status box at the bottom of the window.
- Provides access to other debugger windows from the menu.

The debugger windows

The debugger has the these windows:

- Breakpoints window – Displays the list of current breakpoints.
- Calls window – Displays the current call stack.
- Classes window – Displays a list of classes currently loaded in the Java VM. In addition, this window displays a list of methods for the currently selected class and a list of static variables for the currently selected class. In this window you can set breakpoints on entry to a method or when a static variable is written.
- Connection window – The Connection window is shown when the debugger is started. You can display it again if you wish to disconnect from the database.
- Exceptions window – You can set a particular exception on which to break, or choose to break on all exceptions.
- Inspection window – Displays current static variables, and allows you to modify them. You can also inspect the value of a Java expression, such as the following:
 - Local variables

- Static variables
- Expressions using the dot operator
- Expressions using subscripts []
- Expressions using parentheses, arithmetic, or logical operators.

For example, the following expressions could be used:

```
x[i].field
q + 1
i == 7
(i + 1) * 3
```

- Locals window – Displays current local variables, and allows you to modify them.
- Status window – Displays messages describing the execution state of the Java VM.

Options

The complete set of options for stepping through source code are displayed on the Run menu. They include the following:

Function	Shortcut key	Description
Run	F5	Continue running until the next breakpoint, until the Stop item is selected, or until execution finishes.
Step Over	F7 or Space	Step to the next line in the current method. If the line steps into a different method, step over the method, not into it. Also, step over any breakpoints within methods that are stepped over.
Step Into	F8 or i	Step to the next line of code. If the line steps into a different method, step into the method.

Function	Shortcut key	Description
Step Out	F11	Complete the current method, and break at the next line of the calling method.
Stop		Break execution.
Run to Selected	F6	Run until the currently selected line is executed and then break.
Home	F4	Select the line where the execution is broken.

Setting breakpoints

When you set a breakpoint in the debugger, the Java VM stops execution at that breakpoint. Once execution is stopped, you can inspect and modify the values of variables and other expressions to better understand the state of the program. You can then trace through execution step by step to identify problems.

Setting breakpoints in the proper places is a key to efficiently pinpointing the problem execution steps.

The Java debugger allows you to set breakpoints not only on a line of code, but on many other conditions. This section describes how to set breakpoints using different conditions.

Breaking on a line number

When you break on a particular line of code, execution stops whenever that line of code is executed.

To set a breakpoint on a particular line:

- In the Source window, select the line and press F9.
- Alternatively, you can double-click a line.

When a breakpoint is set on a line number, the breakpoint is shown in the Source window by an asterisk in the left column. If the Breakpoints window is open, the method and line number is displayed in the list of breakpoints.

You can toggle the breakpoint on and off by repeatedly double-clicking or pressing F9.

Breaking on a static method

When you break on a method, the break point is set on the first line of code in the method that contains an executable statement.

To set a breakpoint on a static method:

- 1 From the Source window, choose Break→New. The Break At window is displayed.
- 2 Enter the name of a method in which you wish execution to stop. For example:

```
JDBCExamples.selector
```

stops execution whenever the `JDBCExamples.selector()` method is entered.

When a breakpoint is set on a method, the breakpoint is shown in the Source window by an asterisk in the left column of the line where the breakpoint actually occurs. If the Breakpoints window is open, the method is displayed in the list of breakpoints.

Using counts with breakpoints

If you set a breakpoint on a line that is in a loop, or in a method that is frequently invoked, you may find that the line is executed many times before the condition you are really interested in takes place. The debugger allows you to associate a count with a breakpoint, so that execution stops only when the line is executed a set number of times.

To associate a count with a breakpoint:

- 1 From the Source window, select Break→Display. The Breakpoints window is displayed.
- 2 In the Breakpoints window, click a breakpoint to select it.
- 3 Select Break→Count. A window is displayed with a field for entering a number of iterations. Enter an integer value. The execution will stop when the line has been executed the specified number of times.

Using conditions with breakpoints

The debugger allows you to associate a condition with a breakpoint, so that execution stops only when the line is executed and the condition is met.

To associate a condition with a breakpoint:

- 1 From the Source window, select Break→Display. The Breakpoints window is displayed.
- 2 In the Breakpoints window, click a breakpoint to select it.
- 3 Select Break→Condition. A window is displayed with a field for entering an expression. The execution will stop when the condition is true.

The expressions used here are the same as those that can be used in the Inspection window, and include the following:

- Local variables
- Static variables
- Expressions using the dot operator
- Expressions using subscripts []
- Expressions using parentheses, arithmetic, or logical operators.

Breaking when execution is not interrupted

With a single exception, breakpoints can only be set when program execution is interrupted. If you clear all breakpoints, and run the program you are debugging to completion, you can no longer set a breakpoint on a line or at the start of a method. Also, if a program is running in a loop, execution is continuing and is not interrupted.

To debug your program under either of these conditions, select Run→Stop from the Source window. This stops execution at the next line of Java code that is executed. You can then set breakpoints at other points in the code.

Disconnecting from the database

When the program has run to completion, or at anytime during debugging, you can disconnect from the database from the Connect window. Then, exit the Source window and reconnect to the database after the debug program terminates.

A debugging tutorial

This section takes you through a simple debugging session.

Before you begin

The source code for the class used in this tutorial is located in the directory contained in the zip file

`$SYBASE/$SYBASE_ASE/sample/JavaXml/JavaXml.zip`. See `/JavaXml/Java/Java-Sql-examples` in the unzipped directory.

Before you run the debugger, compile the source code using the `javac` command with the `-g` option.

See “Creating Java classes and JARs” on page 14 for complete instructions for compiling and installing Java classes in the database.

Start the Java debugger and connect to the database

You can start the debugger and connect to the database using a script, command line options, or Sybase Central. In this tutorial, we use *jdebug* to start the debugger. You can use any database.

Follow these steps:

- 1 Start Adaptive Server.
- 2 If Java queries have not yet been executed on your server, run any Java query to initialize the Java subsystem and start a Java VM.
- 3 Run the `$SYBASE/$SYBASE_ASE/debugger/jdebug` script. *jdebug* prompts you for these parameters:
 - a Machine name of the Adaptive Server
 - b Port number for the database
 - c Your login name
 - d Your password
 - e An alternate path to *Debug.jar* if its location is not in your CLASSPATH

Once the connection is established, the debugger window displays a list of available Java VMs or “Waiting for a VM.”

Attach to a Java VM

To attach to a Java VM from your user session:

- 1 With the debugger running, connect to the sample database from isql as the sa:

```
$SYBASE/$SYBASE_OCS/bin/isql -Usa -P
```

Note You cannot start Java execution from the debugger. To start a Java VM you must carry out a Java operation from another connection using the same user name.

- 2 Execute Java code using the following statements:

```
select JDBCExamples.serverMain('createtable')
select JDBCExamples.serverMain('insert')
select JDBCExamples.serverMain('select')
```

The Sybase Java VM starts in order to retrieve the Java objects from the table. The debugger immediately stops execution of the Java code.

The debugger Connection window displays the Java VMs belonging to the user in this format:

```
VM#: "login_name, spid:spid#"
```

- 3 In the debugger Connection window, click the Java VM you want and then click Attach to VM. The debugger attaches to the Java VM and the Source window appears. The Connection window disappears.

Next, enable the Source window to show the source code for the method. The source code is available on disk.

Load source code into the debugger

The debugger looks for source code files. You need to make the `$SYBASE/$SYBASE_ASE/sample/JavaSql/manual-examples/` subdirectory available to the debugger, so that the debugger can find source code for the class currently executing in the database.

To add a source code location to the debugger:

- 1 From the Source window, select File→Source Path. The Source Path window displays.

- 2 From the Source Path window, select Path→Add. Enter the following location into the text box:

```
$SYBASE/$SYBASE_ASE/sample/JavaSql/  
manual-examples/
```

The source code for the `JDBCExamples` class displays in the window, with the first line of the `Query` method `serverMain()` highlighted. The Java debugger has stopped execution of the code at this point.

You can now close the Source Path window.

Step through source code

You can step through source code in the Java debugger in several ways. In this section we illustrate the different ways you can step through code using the `serverMain()` method.

When execution pauses at a line until you provide further instructions, we say that the execution **breaks** at the line. The line is a **breakpoint**. Stepping through code is a matter of setting explicit or implicit breakpoints in the code, and executing code to that breakpoint.

Following the previous section, the debugger should have stopped execution of `JDBCExamples.serverMain()` at the first statement:

Examples

Here are some steps you can try:

- 1 Stepping into a function – press F7 to step to the next line in the current method.
- 2 Press F8 to step into the function `doAction()` in line 99.
- 3 Run to a selected line. You are now in function `doAction()`. Click on line 155 and press F6 to run to that line and break:

```
String workString = "Action(" + action + ")";
```

- 4 Set a breakpoint and execute to it – select line 179 and press F9 to set a breakpoint on that line when running `isql select`
`JDBCExamples.serverMain('select')`:

```
workString + = selector(con);
```

Press F5 to execute to that line.

- 5 Experiment – try different methods of stepping through the code. End with F5 to complete the execution.

When you have completed the execution, the Interactive Data window displays:

```
Action(select) - Row with id = 1: name(Joe Smith)
```

Inspecting and modifying variables

You can inspect the values of both local variables (declared in a method) and class static variables in the debugger.

Inspecting local variables

You can inspect the values of local variables in a method as you step through the code, to better understand what is happening.

To inspect and change the value of a variable:

- 1 Set a breakpoint at the first line of the `selecter()` method from the Breakpoint window. This line is:

```
String sql = "select name, home from xmp where  
id=?";
```

- 2 In Interactive, enter the following statement again to execute the method:

```
select JDBCExamples.serverMain('select')
```

The query executes only as far as the breakpoint.

- 3 Press F7 to step to the next line. The variable has now been declared and initialized.
- 4 From the Source window, select Window→Locals. The Local window appears.

The Locals window shows that there are several local variables. The `sql` variable has a value of zero. All others are listed as not in scope, which means they are not yet initialized.

You must add the variables to the list in the Inspect window.

- 5 In the Source window, press F7 repeatedly to step through the code. As you do so, the values of the variables appear in the Locals window.

If a local variable is not a simple integer or other quantity, then as soon as it is set a + sign appears next to it. This means the local variable has fields that have values. You can expand a local variable by double-clicking the + sign or setting the cursor on the line and pressing **Enter**.

- 6 Complete the execution of the query to finish this exercise.

Modifying local variables

You can also modify values of variables from the Locals window.

To modify a local variable:

- 1 In the debugger Source window, set a breakpoint at the following line in the `selector()` method of the `serverMain` class:

```
String sql = "select name, home from xmp where  
id=?";
```
- 2 Step past this line in the execution.
- 3 Open the Locals window. Select the *id* variable, and select **Local→Modify**. Alternatively, you can set the cursor on the line and press **Enter**.
- 4 Enter a value of 2 in the text box, and click **OK** to confirm the new value. The *id* variable is set to 2 in the Locals window.
- 5 From the Source window, press **F5** to complete execution of the query. In the **Interactive Data** window, an error message displays indicating that no rows were found.

Inspecting static variables

You can also inspect the values of class-level variables (static variables).

To inspect a static variable:

- 1 From the debugger Source window, select **Window→Classes**. The **Classes** window is displayed.
- 2 Select a class in the left box. The methods and static variables of the class are displayed in the boxes on the right.
- 3 Select **Static→Inspect**. The **Inspect** window is displayed. It lists the variables available for inspection.

Adaptive Server 12.5 supports `java.net`, a package that allows you to create networking applications and access different kinds of external servers.

Topic	Page
Overview	123
<code>java.net</code> classes	124
Setting up <code>java.net</code>	124
Example usage	125
User notes	130

Adaptive Server `java.net` is compliant with the Java 1.2 API.

Overview

Support for `java.net` in the Adaptive Server allows you to create client-side Java networking applications within the server. You can create a network Java client application in the Adaptive Server that connects to any server, which in effect enables Adaptive Server to function as a client to external servers. See “Example usage” on page 125.

You can use `java.net` for many purposes:

- Download documents from any URL address on the Internet.
- Send e-mail messages from inside the server.
- Connect to an external server to save a document and perform file functions: saving a document, editing a document, and so forth.
- Access documents using XML.

java.net classes

Table 7-1 shows the java.net classes Sybase supports.

Table 7-1: Supported java.net classes

Class	Supported	Special circumstances
InetAddress	Yes	None
Socket	Yes	Does not support deprecated constructor "Socket (string host, int port, boolean stream)" when stream = false
URL	Yes	No file URL
URLConnection	Yes	None
URLConnection	Yes	No file URL
URLDecoder	Yes	None
URLEncoder	Yes	None
DatagramPacket	No	
DatagramSocket	No	
MulticastSocket	No	
ServerSocket	No	

You can use any of the supported classes in java.net to write Adaptive Server client applications.

Setting up java.net

The following steps enable java.net.

❖ **enabling java.net**

- 1 Enable Java Virtual Machine (VM).

```
sp_configure "enable java", 1
```

- 2 Specify the number of sockets you want to open (the default is 0). The number of sockets configuration parameter is dynamic; you need not restart Adaptive Server if you change the configuration option. For example, to open 10 sockets, enter

```
sp_configure "number of java sockets", 10
```

- 3 Adjust the amount of memory available for the Java VM. Since you may be streaming large text documents in and out, you may need to increase the amount of memory available to the Java VM. The parameters you may need to adjust are:

- size of global fixed heap
- size of process object heap
- size of shared class heap

For more information on these parameters, see Chapter 5, “Configuration Parameters,” in the Sybase *System Administration Guide*.

Example usage

This section provides examples for using both socket classes and the URL class. You can:

- Access an external document with XQL, using the URL class
- Save text out of Adaptive Server
- Use the MailTo class URL to mail a document

Using socket classes

Socket classes allow you to do more sophisticated network transfers than you can achieve using URL classes. The `Socket` class allows you to connect to specified port on any specified network host, and use the `InputStream` and `OutputStream` classes to read and write the data.

Saving text out of Adaptive Server

This example describes how to set up a client application in Adaptive Server. Adaptive Server version 12.5 and later does not support direct access to a file; this example is a workaround for this limitation.

You can write your own external server, which performs file operations, and connect to this new server from the Adaptive Server, using a socket created from a `Socket` class.

In the basic roles of client and server, the client connects to the server and streams the text, while the server receives the stream and streams it to a file.

This example shows how you can install a Java application in Adaptive Server, using `java.net`. This application acts as a client to an external server.

❖ **The client process:**

- 1 Receives an `InputStream`.
- 2 Creates a socket using the `Socket` class to connect to the server.
- 3 Creates an `OutputStream` on the socket.
- 4 Reads the `InputStream` and writes it to the `OutputStream`:

```
import java.io.*;
import java.net.*;
public class TestStream2File {
    public static void writeOut(InputStream fin)throws Exception
    {
        Socket socket = new Socket("localhost", 1718);
        OutputStream fout =
newBufferedOutputStream(socket.getOutputStream());
        byte[] buffer = new byte[10];
        int bytes_read;
        while ((bytes_read = fin.read(buffer)) != -1) {
            fout.write(buffer, 0, bytes_read);
        }
        fout.close();
    }
}
```

Compile this program.

❖ **The server process:**

- 1 Creates a server socket, using the `SocketServer` class, to listen on a port.
- 2 Uses the server socket to obtain a socket connection.
- 3 Receives an `InputStream`.

4 Reads the InputStream and writes it to a FileOutputStream.

Note In this example, the server does not use threads, and therefore it can receive a connection from only one client at a time.

```
import java.io.*;
import java.net.*;
public class FileServer {
    public static void main (string[] args) throws IOException{
        Socket client = accept (1718);
        try{
            InputStream in = client.getInputStream ();
            FileOutputStream fout = new
            FileOutputStream("chastity.txt");
            byte[] buffer = new byte [10];
            int bytes_read;
            while (bytes_read = in.read(buffer))!= -1){
                fout.write(buffer, 0, bytes_read);
            }
            fout.close();
        }
        finally {
            client.close ();
        }
    }
}
static Socket accept (int port) throwsIOException {
    System.out.println ("Starting on port " + port);
    ServerSocket server = new ServerSocket (port);
    System.out.println ("Waiting");
    Socket client = server.accept ();
    System.out.println ("Accepted from " + client.getInetAddress ());
    server.close ();
    return client;
}
}
```

Compile this program.

To use this combination of client and server, you must install the client in Adaptive Server and start the external server:

```
witness% java FileServer &
[2] 28980
witness% Starting on port 1718
```

Waiting

Invoke the client from within Adaptive Server.

```
use pubs2
go
select TestStream2File.writeOut(c1) from blurbs
where au_id = "486-29-1786"
go
```

Using the URL class

You can use the URL class to:

- Send an e-mail message.
- Download an HTTP document from a Web server. This document can be a static file or can be dynamically constructed by the Web server.
- Access an external document with XQL

Use the mailto:URL class to mail a document

Mailing a document is a good example of using the URL class. Before you start, your client must connect to a mail server, so that the machine referenced by System Properties (in this case salsa.sybase.com) is running a mail server, such as sendmail.

- 1 Create a URL object.
- 2 Set a URLConnection object.
- 3 Create an OutputStream object from the URL object.
- 4 Write the mail. For example:

```
import java.io.*;
import java.net.*;
public class MailTo {
    public static void sendIt() throws Exception{
        System.getProperty("mail.host", "salsa.sybase.com");
        URL url = new URL("mailto:name@sybase.com");
        URLConnection conn = url.openConnection();
        PrintStream out = new PrintStream(conn.getOutputStream(),
true);
        out.print ("From: kennys@sybase.com"+"\\r\\n");
        out.print ("Subject: Works Great!"+"\\r\\n");
```

```
        out.print ("Thanks for the example - it works great!"+"\\r\\n");
        out.close();
        System.out.println("Message Sent");
    }
}
```

5 Install mailto:URL for sending e-mail from within the database:

```
select MailTo.sendIt()
Message Sent!
```

A connection to a server is required for these actions.

Obtaining an HTTP document

Another way to use the URL class is to download a document from an HTTP URL. When you start the client connects to a Web server. In the client code, you:

- Create a URL object.
- Create an `InputStream` object from the URL object.
- Use `read` on the `InputStream` object to read in the document.

The following code sample works by:

- Reading the entire document into Adaptive Server memory.
- Creating a new `InputStream` on the document in Adaptive Server memory.

```
import java.io.*;
import java.net.*;
public class URLprocess {
    public static InputStream readURL()
        throws Exception {
        URL u = newURL("http://www.xxxx.com");
        InputStream in = u.openStream();
        //This is the same as creating URLConnection, then
        //calling getInputStream(). In ASE you need to read
        //the entire document into memory, then create an
        //InputStream on the in-memory copy.
        int n=0, off;
        byte b[]=new byte[50000];
        for (off=0; (off<b.length-1); off+=n) {
            n=(in.read(b, off, 512) != -1) ? 512 : b.length-off;
        }
        System.out.println("Number of bytes read : " + off);
        in.close();
        ByteArrayInputStream test =
```

```
        new ByteArrayInputStream(b, 0, off);
    return (InputStream) test;
    }
}
```

After you create the new `InputStream` class, you can install this class and use it to read a text file into the database, inserting data into a table, as in the following example.

```
create table t (cl text)
go

insert into t values (URLprocess.readURL())
go
Number of bytes read :40867

select datalength(cl) from t
go
-----
         40867
```

User notes

Certain aspects of `java.net` require caution:

- Most objects associated with `java.net` are not serializable, which means that you cannot insert them into tables.
- You might encounter the exception “Too many open files,” when you have opened only a few. Check `Number of Java Sockets` configuration parameter.
- Most of the I/O-related functions use buffered I/O, which means that you might need to flush your data explicitly. The `PrintWriter` class is an example of a class in which the data is not automatically flushed.

This chapter presents information on several reference topics.

Topic	Page
JDK requirement for Java classes in the server	131
Assignments	132
Allowed conversions	133
Transferring Java-SQL objects to clients	134
Supported Java API packages, classes, and methods	134
Invoking SQL from Java	137
Transact-SQL commands from Java methods	138
Datatype mapping between Java and SQL	142
Java-SQL identifiers	144
Java-SQL class and package names	145
Java-SQL column declarations	146
Java-SQL variable declarations	147
Java-SQL column references	147
Java-SQL member references	148
Java-SQL method calls	149

JDK requirement for Java classes in the server

Java classes that you install and use in the server *must* be compiled with JDK 1.2.2. If you compile a class with a later JDK, you will be able to install it in the server using the `installjava` utility, but you will receive a `java.lang.ClassFormatError` exception when you attempt to use the class in Adaptive Server.

Assignments

This section defines the rules for assignment between SQL data items whose datatypes are Java-SQL classes.

Each assignment transfers a *source instance* to a *target data item*:

- For an `insert` statement specifying a table that has a Java-SQL column, refer to the Java-SQL column as the target data item and the insert value as the source instance.
- For an `update` statement that updates a Java-SQL column, refer to the Java-SQL column as the target data item and the update value as the source instance.
- For a `select` or `fetch` statement that assigns to a variable or parameter, refer to the variable or parameter as the target data item and the retrieved value as the source instance.

Note If the source is a variable or parameter, then it is a reference to an object in the Java VM. If the source is a column reference, which contains a serialization, then the rules for column references (see [Java-SQL column references on page 147](#)) yield a reference to an object in the Java VM. Thus, the source is a reference to an object in the Java VM.

Assignment rules at compile-time

- 1 Define SC and TC as compile-time class names of the source and target. Define SC_T and TC_T as classes named SC and DT in the database associated with the target. Similarly, define SC_S and TC_S as classes named SC and DT in the database associated with the source.
- 2 SC_T must be the same as TC_T or a subclass of TC_T.

Assignment rules at runtime

Assume that DT_SC is the same as DT_TC or its subclass.

- Define RSC as the runtime class name of the source value. Define RSC_S as the class named RSC in the database associated with the source. Define RSC_T as the name of a class RSC_T installed in the database associated with the target. If there is no class RSC_T, then an exception is raised. If RSC_T is neither the same as TC_T nor a subclass of TC_T, then an exception is raised.
- If the databases associated with the source and target are not the same database, then the source object is serialized by its current class, RSC_S, and that serialization is deserialized by the class RSC_T that it will be associated with in the database associated with the target.
- If the target is a SQL variable or parameter, then the source is copied by reference to the target.
- If the target is a Java-SQL column, then the source is serialized, and that serialization is deep copied to the target.

Allowed conversions

You can use `convert` to change the expression datatype in these ways:

- Convert Java types where the Java datatype is a Java object type to the SQL datatype shown in “Datatype mapping between Java and SQL” on page 142. The action of the `convert` function is the mapping implied by the Java-SQL mapping.
- Convert SQL datatypes to Java types shown in “Datatype mapping between Java and SQL” on page 142. The action of the `convert` function is the mapping implied by the SQL-Java mapping.
- Convert any Java-SQL class installed in the SQL system to any other Java-SQL class installed in the SQL system if the compile-time datatype of the expression (source class) is a subclass or superclass of the target class. Otherwise, an exception is raised.

The result of the conversion is associated with the current database.

See “Using the SQL `convert` function for Java subtypes,” for a discussion of the use of the `convert` function for Java subtypes.

Transferring Java-SQL objects to clients

When a value whose datatype is a Java-SQL object type is transferred from Adaptive Server to a client, the data conversion of the object depends on the client type:

- If the client is an isql client, the `toString()` or similar method of the object is invoked and the result is truncated to `varchar`, which is transferred to the client.

Note The number of bytes transferred to the client is dependent on the value of the `@@stringsize` global variable. The default value is 50 bytes. See “Representing Java instances” on page 30 for more information.

- If the client is a Java client that uses `jdbcConnect` 4.0 or later, the server transmits the object serialization to the client. This serialization is seamlessly deserialized by `jdbcConnect` to yield a copy of the object.
- If the client is a `dbclient`:
 - If the object is a column declared as `in row`, the serialized value contained in the column is transferred to the client as a `varbinary` value of length determined by the size of the column.
 - Otherwise, the serialized value of the object (the result of the `writeObject` method of the object) is transferred to the client as an image value.

Supported Java API packages, classes, and methods

Adaptive Server supports many but not all classes and methods in the Java API. In addition, Adaptive Server may impose security restrictions and implementation limitations. For example, Adaptive Server does not support all of the thread creation and manipulation facilities of `java.lang.Thread`.

The supported packages are installed with Adaptive Server and are always available. They cannot be installed by the user.

Note Java in Adaptive Server does not support the Java Native Interface (JNI).

This section lists:

- Supported Java packages and classes
- Unsupported Java packages
- Unsupported java.sql methods

Supported Java packages and classes

- java.io
 - Externalizable
 - DataInput
 - DataOutput
 - ObjectInputStream
 - ObjectOutputStream
 - Serializable
- java.lang – see “Unsupported java.sql methods and interfaces” on page 136 for a list of the unsupported classes in java.lang.
- java.math
- java.net – see Chapter 7, “Network Access Using java.net”
- java.sql – see “Unsupported java.sql methods and interfaces” on page 136 for a list of the unsupported methods and interfaces in java.sql.
- java.text
- java.util
- java.util.zip

Unsupported Java packages, classes, and methods

- java.applet
- java.awt
- java.awt.datatransfer
- java.awt.event
- java.awt.image

- `java.awt.peer`
- `java.beans`
- `java.lang.ref`
- `java.lang.Thread`
- `java.lang.ThreadGroup`
- `java.rmi`
- `java.rmi.dgc`
- `java.rmi.registry`
- `java.rmi.server`
- `java.security`
- `java.security.acl`
- `java.security.interfaces`

Unsupported *java.sql* methods and interfaces

- `Connection.commit()`
- `Connection.getMetaData()`
- `Connection.nativeSQL()`
- `Connection.rollback()`
- `Connection.setAutoCommit()`
- `Connection.setCatalog()`
- `Connection.setReadOnly()`
- `Connection.setTransactionIsolation()`
- `DatabaseMetaData.*` – `DatabaseMetaData` is supported except for these methods:
 - `deletesAreDetected()`
 - `getUDTs()`
 - `insertsAreDetected()`
 - `updatesAreDetected()`

- `othersDeletesAreVisible()`
- `othersInsertsAreVisible()`
- `othersUpdatesAreVisible()`
- `ownDeletesAreVisible()`
- `ownInsertsAreVisible()`
- `ownUpdatesAreVisible()`
- `PreparedStatement.setAsciiStream()`
- `PreparedStatement.setUnicodeStream()`
- `PreparedStatement.setBinaryStream()`
- `ResultSetMetaData.getCatalogName()`
- `ResultSetMetaData.getSchemaName()`
- `ResultSetMetaData.getTableName()`
- `ResultSetMetaData.isCaseSensitive()`
- `ResultSetMetaData.isReadOnly()`
- `ResultSetMetaData.isSearchable()`
- `ResultSetMetaData.isWritable()`
- `Statement.getMaxFieldSize()`
- `Statement.setMaxFieldSize()`
- `Statement.setCursorName()`
- `Statement.setEscapeProcessing()`
- `Statement.getQueryTimeout()`
- `Statement.setQueryTimeoutt()`

Invoking SQL from Java

Adaptive Server supplies a native JDBC driver, `java.sql`, that implements JDBC 1.1 and 1.2 specifications, and is compliant with version 2.0. `java.sql` enables Java methods executing in Adaptive Server to perform SQL operations.

Special considerations

`java.sql.DriverManager.getConnection()` accepts these URLs:

- `null`
- `""` (the null string)
- `jdbc:default:connection`

When invoking SQL from Java some restrictions apply:

- A SQL query that is performing update actions (update, insert, or delete) cannot use the facilities of `java.sql` to invoke other SQL operations that also perform update actions.
- Triggers that are fired by SQL using the facilities of `java.sql` cannot generate result sets.
- `java.sql` cannot be used to execute extended stored procedures or remote stored procedures.

Transact-SQL commands from Java methods

You can use certain Transact-SQL commands in Java methods called within the SQL system. Table 8-1 lists Transact-SQL commands and whether or not you can use them in Java methods. You can find further information on most of these commands in the *Sybase Adaptive Server Enterprise Reference Manual*.

Table 8-1: Support status of Transact-SQL commands

Command	Status
alter database	Not supported.
alter role	Not supported.
alter table	Supported.
begin ... end	Supported.
begin transaction	Not supported.
break	Supported.
case	Supported.
checkpoint	Not supported.
commit	Not supported.
compute	Not supported.

Command	Status
connect - disconnect	Not supported.
continue	Supported.
create database	Not supported.
create default	Not supported.
create existing table	Not supported.
create function	Supported.
create index	Not supported.
create procedure	Not supported.
create role	Not supported.
create rule	Not supported.
create schema	Not supported.
create table	Supported.
create trigger	Not supported.
create view	Not supported.
cursors	Not supported. Only “server cursors” are supported, that is, cursors that are declared and used within a stored procedure.
dbcc	Not supported.
declare	Supported.
disk init	Not supported.
disk mirror	Not supported.
disk refit	Not supported.
disk reinit	Not supported.
disk remirror	Not supported.
disk unmirror	Not supported.
drop database	Not supported.
drop default	Not supported.
drop function	Supported.
drop index	Not supported.
drop procedure	Not supported.
drop role	Not supported.
drop rule	Not supported.
drop table	Supported.
drop trigger	Not supported.
drop view	Not supported.

Command	Status
dump database	Not supported.
dump transaction	Not supported.
execute	Supported.
goto	Supported.
grant	Not supported.
group by and having clauses	Supported.
if...else	Supported.
insert table	Supported.
kill	Not supported.
load database	Not supported.
load transaction	Not supported.
online database	Not supported.
order by Clause	Supported.
prepare transaction	Not supported.
print	Not supported.
raiserror	Supported.
readtext	Not supported.
return	Supported.
revoke	Not supported.
rollback trigger	Not supported.
rollback	Not supported.
save transaction	Not supported.
set	See Table 12-2 for set options.
setuser	Not supported.
shutdown	Not supported.
truncate table	Supported.
union Operator	Supported.
update statistics	Not supported.
update	Supported.
use	Not supported.
waitfor	Supported.
where Clause	Supported.
while	Supported.
writetext	Not supported.

Table 8-2 lists set command options and whether or not you can use them in Java methods.

Table 8-2: Support status of set command options

set command option	Status
ansinull	Supported.
ansi_permissions	Supported.
arithabort	Supported.
arithignore	Supported.
chained	Not supported. See Note 1.
char_convert	Not supported.
cis_rpc_handling	Not supported
close on endtran	Not supported
cursor rows	Not supported
datefirst	Supported
dateformat	Supported
fipsflagger	Not supported
flushmessage	Not supported
forceplan	Supported
identity_insert	Supported
language	Not supported
lock	Supported
nocount	Supported
noexec	Not supported
offsets	Not supported
or_strategy	Supported
parallel_degree	Supported. See Note 2.
parseonly	Not supported
prefetch	Supported
process_limit_action	Supported. See Note 2.
procid	Not supported
proxy	Not supported
quoted_identifier	Supported
replication	Not supported
role	Not supported
rowcount	Supported
scan_parallel_degree	Supported. See Note2.
self_recursion	Supported

set command option	Status
session_authorization	Not supported
showplan	Supported
sort_resources	Not supported
statistics io	Not supported
statistics subquerycache	Not supported
statistics time	Not supported
string_rtruncation	Supported
stringsize	Supported
table count	Supported
textsize	Not supported
transaction iso level	Not supported. See Note 1.
transactional_rpc	Not supported

Note (1) set commands with options chained or transaction isolation level are allowed only if the setting that they specify is already in effect. That is, this kind of set command is allowed if it has no affect. This is done to support common coding practises in stored procedures.

Note (2) set commands pertaining to parallel degree are allowed but have no affect. This supports the use of stored procedures that set the parallel degree for other contexts.

Datatype mapping between Java and SQL

Adaptive Server maps SQL datatypes to Java types (SQL-Java datatype mapping) and Java scalar types to SQL datatypes (Java-SQL datatype mapping). Table 8-3 shows SQL-Java datatype mapping.

Table 8-3: Mapping SQL datatypes to Java types

SQL type	Java type
char	String
varchar	String
nchar	String
nvarchar	String
unichar	String
univarchar	String
unitext	String
text	String
numeric	java.math.BigDecimal
decimal	java.math.BigDecimal
money	java.math.BigDecimal
smallmoney	Java.math.BigDecimal
bit	boolean
tinyint	byte
smallint	short
integer	int
bigint	long
unsigned smallint	int
unsigned int	long
unsigned bigint	java.math.BigInteger
bigint	java.math.BigInteger
real	float
float	double
double precision	double
binary	byte[]
varbinary	byte[]
image	java.io.InputStream
datetime	java.sql.Timestamp
smalldatetime	java.sql.Timestamp
date	java.sql.Date
time	java.sql.Time

Note The mapping of unsigned bigint to double is an approximation; it will not provide exact values. For exact values, convert the unsigned bigint value to a string value when passing it to a Java method.

Table 8-4 shows Java-SQL datatype mapping.

Table 8-4: Mapping Java scalar types to SQL datatypes

Java scalar type	SQL type
boolean	bit
byte	tinyint
short	smallint
int	integer
long	bigint
float	real
double	double

Java-SQL identifiers

Description	Java-SQL identifiers are a subset of Java identifiers that can be referenced in SQL.
Syntax	<code>java_sql_identifier ::= alphabetic character underscore (_) symbol [alphabetic character arabic numeral underscore (_) symbol dollar (\$) symbol]</code>
Usage	<ul style="list-style-type: none">• Java-SQL identifiers can be a maximum of 255 bytes in length if they are surrounded by quotation marks. Otherwise, they must be 30 bytes or fewer.• The first character of the identifier must be either an alphabetic character (uppercase or lowercase) or the underscore (_) symbol. Subsequent characters can include alphabetic characters (uppercase or lowercase), numbers, the dollar (\$) symbol, or the underscore (_) symbol.• Java-SQL identifiers are always case sensitive.

Delimited Identifiers

- Delimited identifiers are object names enclosed in double quotes. Using delimited identifiers for Java-SQL identifiers allows you to avoid certain restrictions on the names of Java-SQL identifiers.

Note You can use double quotes with Java-SQL identifiers whether the `quoted_identifier` option is on or off.

- Delimited identifiers allow you to use SQL reserved words for packages, classes, methods, and so on. Each time you use the delimited identifier in a statement, you must enclose it in double quotes. For example:

```
create table t1
(c1 char(12)
c2 p1."select".p2."jar")
```

- Double quotes surround only individual Java-SQL identifiers, not the fully qualified name.

See also

For additional information about identifiers, see Chapter 5, “Transact-SQL Topics,” in the *Reference Manual*.

Java-SQL class and package names

Description	To reference a Java-SQL class or package, use the following syntax:
Syntax	<pre>java_sql_class_name ::= [java_sql_package_name.]java_sql_identifier java_sql_package_name ::= [java_sql_package_name.]java_sql_identifier</pre>
Parameters	<p><i>java_sql_class_name</i> The fully qualified name of a Java-SQL class in the current database.</p> <p><i>java_sql_package_name</i> The fully qualified name of a Java-SQL package in the current database.</p> <p><i>java_sql_identifier</i> See Java-SQL identifiers.</p>
Usage	<p>For Java-SQL class names:</p> <ul style="list-style-type: none"> A class name reference always refers to a class in the current database.

- If you specify a Java-SQL class name without referencing the package name, only one Java-SQL class of that name must exist in the current database, and its package must be the default (anonymous) package.
- If a SQL user-defined datatype and a Java-SQL class possess the same sequence of identifiers, Adaptive Server uses the SQL user-defined datatype name and ignores the Java-SQL class name

For Java-SQL package names:

- If you specify a Java-SQL subpackage name, you must reference the subpackage name with its package name:
java_sql_package_name.java_sql_subpackage_name
- Use Java-SQL package names only as qualifiers for class names or subpackage names and to delete packages from the database using the `remove java` command.

Java-SQL column declarations

Description	To declare a Java-SQL column when you create or alter a table, use the following syntax:
Syntax	<code>java_sql_column ::= column_name java_sql_class_name</code>
Parameters	<p><i>java_sql_column</i> Specifies the syntax of Java-SQL column declarations.</p> <p><i>column_name</i> The name of the Java-SQL column.</p> <p><i>java_sql_class_name</i> The name of a Java-SQL class in the current database. This is the “declared class” of the column.</p>
Usage	<ul style="list-style-type: none"> • The declared class must implement either the <code>Serializable</code> or <code>Externalizable</code> interface. • A Java-SQL column is always associated with the current database. • A Java-SQL column cannot be specified as: <ul style="list-style-type: none"> • <code>not null</code> • <code>unique</code> • A primary key

See also You use a Java-SQL column declaration only when you create or alter a table.
 See the create table and alter table information in the *Reference Manual*.

Java-SQL variable declarations

Description	Use Java-SQL variable declarations to declare variables and stored procedure parameters for datatypes that are Java-SQL classes.
Syntax	<i>java_sql_variable</i> ::= @ <i>variable_name</i> <i>java_sql_class_name</i> <i>java_sql_parameter</i> ::= @ <i>parameter_name</i> <i>java_sql_class_name</i>
Parameters	<i>java_sql_variable</i> Specifies the syntax of a Java-SQL variable in a SQL stored procedure. <i>java_sql_parameter</i> Specifies the syntax of a Java-SQL parameter in a SQL stored procedure. <i>java_sql_class_name</i> The name of a Java-SQL class in the current database.
Usage	A <i>java_sql_variable</i> or <i>java_sql_parameter</i> is always associated with the database containing the stored procedure.
See also	Refer to the <i>Reference Manual</i> for more information about variable declarations.

Java-SQL column references

Description	To reference a Java-SQL column, use the following syntax:
Syntax	<i>column_reference</i> ::= [[<i>database_name</i> .] <i>owner</i> .] <i>table_name</i> .] <i>column_name</i> <i>database_name</i> .. <i>table_name</i> . <i>column_name</i>
Parameters	<i>column_reference</i> A reference to a column whose datatype is a Java-SQL class.
Usage	<ul style="list-style-type: none"> • If the value of the column is null, then the column reference is also null. • If the value of the column is a Java serialization, S, and the name of its class is CS, then:

- If the class CS does not exist in the current database or if CS is not the name of a class in the database associated with the serialization, then an exception is raised.

Note The database associated with the serialization is normally the database that contains the column. Serializations contained in work tables and in temporary tables created with “insert into #tempdb” are, however, associated with the database in which the serialization was stored originally.

- The value of the column reference is:

CSC.readObject(S)

where CSC is the column reference. If the expression raises an uncaught Java exception, then an exception is raised.

The expression yields a reference to an object in the Java VM, which is associated with the database associated with the serialization.

Java-SQL member references

Description	References a field or method of a class or class instance.
Syntax	<pre>member_reference ::= class_member_reference instance_member_reference class_member_reference ::= java_sql_class_name.method_name instance_member_reference ::= instance_expression>>member_name instance_expression ::= column_reference variable_name parameter_name method_call member_reference member_name ::= field_name method_name</pre>
Parameters	<p><i>member_reference</i> An expression that describes a field or method of a class or object.</p> <p><i>class_member_reference</i> An expression that describes a static method of a Java-SQL class.</p> <p><i>instance_member_reference</i> An expression that describes a static or dynamic method or field of a Java-SQL class instance.</p>

java_sql_class_name

A fully qualified name of a Java-SQL class in the current database.

instance_expression

An expression whose datatype is a Java-SQL class.

member_name

The name of a field or method of the class or class instance.

Usage

- If a member references a field of a class instance, the instance has a null value, and the Java-SQL member reference is the target of a `fetch`, `select`, or `update` statement, then an exception is raised.

Otherwise, the Java-SQL member reference has the null value.

- The double angle (`>>`) and dot (`.`) qualification take precedence over any operator, such as the addition (`+`) or equal to (`=`) operator, for example:

```
X>>A1>>B1 + X>>A1>>B2
```

In this expression, the addition operation is performed after the members have been referenced.

- The field or method designated by a member reference is associated with the same database as that of its Java-SQL class or instance of its Java-SQL class.

If the Java type of a member reference is one of the Java scalar types (such as boolean, byte, and so on), then the corresponding SQL datatype of the reference is obtained by mapping the Java type to its equivalent SQL type.

If the Java type of a member reference is an object type, then the SQL datatype is the same Java object type or class.

Java-SQL method calls

Description To invoke a Java-SQL method, which returns a single value, use the following syntax:

Syntax

```
method_call ::= member_reference ([parameters])
              | new java_sql_class_name ([parameters])
parameters ::= parameter [(, parameter)...]
parameter ::= expression
```

Parameters

method_call

An invocation of a static method, instance method, or class constructor. A method call can be used in an expression where a non-constant value of the method's datatype is required.

member_reference

A member reference that denotes a method.

parameters

The list of parameters to be passed to the method. If there are no parameters, include empty parentheses.

Usage

Method overloading

- When there are methods with the same name in the same class or instance, the issue is resolved according to Java method overloading rules.

Datatype of method calls

- The datatype of a method call is determined as follows:
 - If a method call specifies *new*, its datatype is that of its Java-SQL class.
 - If a method call specifies a member reference that denotes a type-valued method, then the datatype of the method call is that type.
 - If a method call specifies a member reference that denotes a void static method, then the datatype of the method call is SQL integer.
 - If a method call specifies a member reference that denotes a void instance method of a class, then the datatype of the method call is that of the class.
- To include a parameter in a member reference when the parameter is a Java-SQL instance associated with another database, you must ensure that the class name associated with the Java-SQL instance is included in both databases. Otherwise, an exception is raised.

Runtime results

- The runtime result of a method call is as follows:
 - If a method call specifies a member reference whose runtime value is null (that is, a reference to a member of a null instance), then the result is null.
 - If a method call specifies a member reference that denotes a type-valued method, then the result is the value returned by the method.

- If a method call specifies a member reference that denotes a void static method, then the result is the null value.
- If a method call specifies a member reference that denotes a void instance method of an instance of a class, then the result is a reference to that instance.
- The method call and result of the method call are associated with the same database.
- Adaptive Server does not pass the null value as the value of a parameter to a method whose Java type is scalar.

Glossary

This glossary describes Java and Java-SQL terms used in this book. For a description of Adaptive Server and SQL terms, refer to the *Adaptive Server Glossary*.

assignment	A generic term for the data transfers specified by <code>select</code> , <code>fetch</code> , <code>insert</code> , and <code>update</code> Transact-SQL commands. An assignment sets a source value into a target data item.
associated JAR	If a class/JAR is installed with <code>installjava</code> and the <code>-jar</code> option, then the JAR is retained in the database and the class is linked in the database with the associated JAR. See retained JAR .
bytecode	The compiled form of Java source code that is executed by the Java VM.
class	A class is the basic element of Java programs, containing a set of field declarations and methods. A class is the master copy that determines the behavior and attributes of each instance of that class. class definition is the definition of an active data type, that specifies a legal set of values and defines a set of methods that handle the values. See class instance .
class method	See static method .
class file	A file of type “class” (for example, <i>myclass.class</i>) that contains the compiled bytecode for a Java class. See Java file and Java archive (JAR) .
class instance	Value of the class data type that contains a value for each field of the class and that accepts all methods of the class.
datatype mapping	Conversions between Java and SQL datatypes.
declared class	The declared datatype of a Java-SQL data item. It is either the datatype of the runtime value or a supertype of it.
externalization	An externalization of a Java instance is a byte stream that contains sufficient information for the class to reconstruct the instance. Externalization is defined by the externalizable interface. All Java-SQL classes must be either externalizable or serializable. See serialization .

installed classes	Java classes and methods that have been placed in the Adaptive Server system by the <code>installjava</code> utility.
instance method	A invoked method that references a specific instance of a class.
interface	A named collection of method declarations. A class can implement an interface if the class defines all methods declared in the interface.
Java archive (JAR)	A platform-independent format for collecting classes in a single file.
Java Database Connectivity (JDBC)	A Java-SQL API that is a standard part of the Java Class Libraries that control Java application development. JDBC provides capabilities similar to those of ODBC.
Java datatypes	Java classes, either user-defined or from the JavaSoft API, or Java primitive datatypes, such as boolean, byte, short, and int.
Java Development Kit (JDK)	A toolset from Sun Microsystems that allows you to write and test Java programs from the operating system.
Java file	A file of type “java” (for example, <i>myfile.java</i>) that contains Java source code. See class file and Java archive (JAR) .
Java method signature	The Java datatype of each parameter of a Java method.
Java object	An instance of a Java class that is contained in the storage of the Java VM. Java instances that are referenced in SQL are either values of Java columns or Java objects.
Java-SQL column	A SQL column whose datatype is a Java-SQL class.
Java-SQL class	<p>A public Java class that has been installed in the Adaptive Server system. It consists of a set of variable definitions and methods.</p> <p>A class instance consists of an instance of each of the fields of the class. Class instances are strongly typed by the class name.</p> <p>A subclass is a class that is declared to extend (at most) to one other class. That other class is called the direct superclass of the subclass. A subclass has all of the variables and methods of its direct and indirect superclasses, and may be used interchangeably with them.</p>
Java-SQL datatype mapping	Conversions between Java and SQL datatypes. See “Datatype mapping between Java and SQL” on page 142.
Java-SQL variable	A SQL variable whose datatype is a Java-SQL class.

Java Virtual Machine (Java VM)	The Java interpreter that processes Java in the server. It is invoked by the SQL implementation.
mappable	<p>A Java datatype is mappable if it is either:</p> <ul style="list-style-type: none">• Listed in the first column of Table 8-3 on page 143, or• A public Java-SQL class that is installed in the Adaptive Server system. <p>A SQL datatype is mappable if it is either:</p> <ul style="list-style-type: none">• Listed in the first column of Table 8-4 on page 144, or• A public Java-SQL class that is built-in or installed in the Adaptive Server system. <p>A Java method is mappable if all of its parameter and result datatypes are mappable.</p>
method	<p>A set of instructions, contained in a Java class, for performing a task. A method can be declared static, in which case it is called a class method. Otherwise, it is an instance method. Class methods can be referenced by qualifying the method name with either the class name or the name of an instance of the class. Instance methods are referenced by qualifying the method name with the name of an instance of the class. The method body of an instance method can reference the variables local to that instance.</p>
narrowing conversion	<p>A Java operation for converting a reference to a class instance to a reference to an instance of a subclass of that class. This operation is written in SQL with the <code>convert</code> function. See also widening conversion.</p>
package	<p>A package is a set of related classes. A class either specifies a package or is part of an anonymous default package. A class can use Java <code>import</code> statements to specify other packages whose classes can then be referenced.</p>
procedure	<p>An SQL stored procedure, or a Java method with a <i>void</i> result type.</p>
public	<p>Public fields and methods, as defined in Java.</p>
retained JAR	<p>See associated JAR.</p>
serialization	<p>A serialization of a Java instance is a byte stream containing sufficient information to identify its class and reconstruct the instance. All Java-SQL classes must be either externalizable or serializable. See externalization.</p>
SQL function signature	<p>The SQL datatype of each parameter of a SQLJ function.</p>

SQL-Java datatype mapping	Conversions between Java and SQL datatypes. See “Datatype mapping between Java and SQL” on page 142.
SQL procedure signature	The SQL datatype of each parameter of a SQLJ procedure.
static method	A method invoked without referencing an object. Static methods affect the whole class, not an instance of the class. Also called a class method.
subclass	A class below another class in a hierarchy. It inherits attributes and behavior from classes above it. A subclass may be used interchangeably with its superclasses. The class above the subclass is its direct superclass. See superclass , narrowing conversion , and widening conversion .
superclass	A class above one or more classes in a hierarchy. It passes attributes and behavior to the classes below it. It may not be used interchangeably with its subclasses. See subclass , narrowing conversion , and widening conversion .
synonymous classes	Java-SQL classes that have the same fully qualified name but are installed in different databases.
Unicode	A 16-bit character set defined by ISO 10646 that supports many languages.
variable	In Java, a variable is local to a class, to instances of the class, or to a method. A variable that is declared static is local to the class. Other variables declared in the class are local to instances of the class. Those variables are called fields of the class. A variable declared in a method is local to the method.
visible	A Java class that has been installed in a SQL system is visible in SQL if it is declared public; a field or method of a Java instance is visible in SQL if it is both public and mappable. Visible classes, fields, and methods can be referenced in SQL. Other classes, fields, and methods cannot, including classes that are private, protected, or friendly, and fields and methods that are either private, protected, or friendly, or are not mappable.
well-formed document	In XML, the necessary characteristics of a well-formed document include: all elements with both start and end tags, attribute values in quotes, all elements properly nested.
widening conversion	A Java operation for converting a reference to a class instance to a reference to an instance of a superclass of that class. This operation is written in SQL with the <code>convert</code> function. See also narrowing conversion .

Index

Symbols

`::=` (BNF notation)
 in SQL statements xvii
`,` (comma)
 in SQL statements xvii
`{ }` (curly braces)
 in SQL statements xvii
`()` (parentheses)
 in SQL statements xvii
`[]` (square brackets)
 in SQL statements xvii
`>>` (double angle)
 to qualify Java fields and methods 149
`@` sign 81

A

Adaptive Server
 plug-in 25, 80
additional information
 about Java 9
ADT mappable datatypes 98
alter table
 command 25
 syntax 25
ANSI standards 4
assignment properties
 Java-SQL data items 31
assignments 132
attaching to a Java VM 111

B

Backus Naur Form (BNF) notation xvi, xvii
BNF notation in SQL statements xvi, xvii
brackets. *See* square brackets []
breaking

 on a class method 115
 on a line number 114
 using conditions 115
 using counts 115
 when execution is not interrupted 116
breakpoints 114

C

called on null input parameter 82
case expressions 35, 85
character sets
 Adaptive server plug-in 80
 unicode 25, 34, 80
class names 145
class subtypes 34–36
classes. *See* Java classes
clients
 bcp 134
 isql 134
client-side JDBC 6
column
 declarations 146
 referencing 147
column datatypes, requirements 23
column declarations 146
column references 147
comma (,)
 in SQL statements xvii
command main method 101
commands
 create table 24, 25
 drop function 86
 SQLJ create function 81
 SQLJ create procedure 87
commands, create procedure SQLJ 89
compile-time datatypes 36
compiling Java code 14
configuration parameter, Number of Java Sockets 130

- constructor method 26
- constructors 26, 41
- conventions
 - See also* syntax
 - Java-SQL syntax xvi
 - Transact-SQL syntax xvi
- conversions 133
 - narrowing 35
 - widening 35
- convert function 34, 133
- create procedure (SQLJ) command 87, 89
- create table command, syntax 24, 25
- creating
 - client applications 123
 - network applications, java.net 123
 - tables 24
 - user-defined classes 14
- curly braces ({}) in SQL statements xvii

D

- DatagramPacket, Java class 124
- datatype conversions 133
- datatype mapping 33, 97, 142–144
- datatypes
 - compile-time 36
 - conversions 133
 - Java classes 3
 - method calls 150
 - runtime 36
- Debug.jar, Java file 110
- debugger
 - attaching to a Java VM 111
 - compiling classes for 111
 - disconnecting 116
 - how it works 109
 - location 110
 - options 113
 - requirements for using 109
 - starting 110
 - wait mode 111
- debugger capabilities
 - browse classes 110
 - inspect and break on expressions 110
 - inspect and set variables 110

- set break conditions 110
- set breakpoints 110
- trace execution 110
- debugger windows
 - breakpoints 112
 - calls 112
 - classes 112
 - connection 112
 - exceptions 112
 - inspection 112
 - locals 113
 - source 112
- debugging
 - Java 109–121
- debugging tutorial 117–121
 - attaching to a Java VM 118
 - examples 119
 - inspecting local variables 120
 - inspecting static variables 121
 - inspecting variables 119
 - loading source code 118
 - modifying local variables 121
 - source code 117
 - starting the debugger 117
 - stepping through source code 119
- deleting 26, 96
 - Java objects 26
- delimited identifiers 145
- deterministic parameter 82, 88
- disabling Java 13
- distinct keyword 44
- double angle
 - qualifying Java fields and methods 149
 - to qualify Java fields and methods 27
- downloading
 - installed classes 19
 - installed JARs 19
- drop function command 86
- dynamic result sets parameter 88

E

- email
 - java.net 123
 - messages, sending 123

- enabling Java 13
- enabling java.net, procedure 124
- equality operations 44
- examples
 - for SQLJ routines 77
- exceptions 29
- explicit Java method signatures 99
- external name parameter 88
- external server, writing with java.net 125
- externalization 146
- extractjava utility 19

F

- flushing data explicitly 130

G

- group by clause 44

H

- URLConnection, Java class 124

I

- identifiers 144
 - delimited 145
- implicit Java method signatures 99
- in parameter 90
- InetAddress, Java class 124
- inout parameter 90
- InputStream class 128
- InputStream, Java class 130
- inserting
 - data in a table 128
 - Java objects 26
- installing
 - compressed JARS 15
 - Java classes 15, 18
 - uncompressed JARS 15
- installjava utility 12, 15

- f option 16
- j option 16
- new option 17
- syntax 16
- update option 17
- instance methods 42
- inter-class arguments 50
- invoking
 - Java method, using SQLJ 79
 - Java methods 28, 78
 - Java methods, invoking directly 78
 - Java methods, using SQLJ 78
 - SQL from Java 137, 142

J

- JAR files
 - creating 15
 - installing 15
 - retaining 16
- JARs
 - compressed, installing 15
 - uncompressed, installing 15
- Java API 7
 - accessing from SQL 7
 - supported packages 134–137
 - Sybase support for 8
- Java arrays 90
- Java class datatypes 84
- Java classes
 - as datatypes 3, 23
 - creating 14
 - DatagramPacket 124
 - DatagramSocket 124
 - URLConnection
 - InetAddress
 - InputStream 125, 128
 - installing 15–18
 - MailTo 128
 - MulticastSocket 124
 - OutputStream 125, 128
 - PrintWriter 130
 - referencing other classes 18
 - retained 20
 - runtime 12

- saving in JAR 15
- ServerSocket 124, 126
- Socket
- SQLJ examples 78
- subtypes 34
- supported 8
- updating 17
- URL 128, 129
- URL class, using 126
- URLConnection 124
- URLDecoder 124
- URLEncoder 124
- user-defined 8, 12
- Java code
 - compiling 14
 - writing 14
- Java compiler 111
- Java datatypes
 - ADT mappable 98
 - object mappable 98
 - output mappable 98
 - result-set mappable 98
 - simply mappable 98
- Java Development Kit 5
- Java in the database
 - advantages of 1
 - capabilities 2
 - key features 5
 - preparing for 11–20
 - questions and answers 4
- Java instances, representing 30
- Java method signature 83, 88
- Java methods
 - call by reference 29, 45
 - command main 101
 - exceptions 29
 - instance 42
 - invoking 28, 78
 - static 43
 - type 40, 41
 - void 41
- Java objects 26
- Java operations, invoked from SQL 6
- Java primitive datatypes 84
- Java runtime environment 11
- Java VM 6, 11
- Java VM parameters
 - size of global fixed heap 125
 - size of process object heap 125
 - size of shared class heap 125
- Java, SQL, using together 6
- java.net 124, 125, 126, 130
 - accessing documents using XML, JDBC 123
 - accessing external documents 125
 - cautions 130
 - classes
 - client application, setting up 125
 - client process 126
 - client process procedure 126
 - connecting through JDBC with jconnect 123
 - creating networking applications 123
 - downloading documents 123
 - enabling 124
 - examples 125
 - help 130
 - mailing documents 125
 - objects not serializable 130
 - procedure for enabling 124
 - reference documents 130
 - saving documents 123
 - saving text from Adaptive Server 125
 - sending email messages 123
 - server process 126
 - server process procedure 126
 - writing external server 125
- java.net classes
 - URLConnection 124
 - InetAddress 124
 - See Java classes
 - Socket 124
 - URL 124
 - URLConnection 124
 - URLDecoder 124
 - URLEncoder 124
- java.net, for network access 123
- java.sql 137
- java.sql methods, unsupported 136
- Java-SQL
 - class names 145
 - column declarations 146
 - column references 147
 - columns 31, 45

- creating tables 24
- function results 31
- identifiers 144
- member references 148
- method calls 149
- names 23
- package names 145
- parameters 31, 46
- static variables 46
- transferring objects 134
- transferring objects to clients 133
- unsupported methods 136
- variable declarations 147
- variables 31, 46
- Java-SQL classes
 - in multiple databases 46
 - installing 15–18
- Java-SQL columns
 - storage options 24
- jConnect
 - JDBC 6
- jconnect 123
- JDBC 57–74
 - accessing data 59
 - client-side 6, 58
 - concepts 58
 - connection defaults 59
 - connections 62
 - interface 8
 - JDBCExamples class 60
 - obtaining a connection 62
 - permissions 59
 - server-side 6, 58
 - terminology 58
 - version support 12
- JDBC drivers 12, 137
 - client-side 6, 58
 - jConnect 6
 - server-side 6, 58
- JDBC standard datatype mapping 97
- JDBCExamples class 68–74
 - methods 61–66
 - overview 60

L

- language java parameter 88

M

- mailing a document 125
- MailTo, Java class 128
- mapping datatypes 142–144
- mapping Java and SQL datatypes 97
- member references 148
- method calls 149
 - datatype of 150
- method overloading 100, 150
- methods
 - exceptions 29
 - runtime results 150
 - See also XQL methods
 - SQLJExamples.bestTwoEmps() 78
 - SQLJExamples.correctStates() 78, 89
 - SQLJExamples.job() 78
 - SQLJExamples.region() 78
- modifies sql data parameter 82, 88
- MulticastSocket, Java class 124
- multiple databases 47

N

- names in Java-SQL 23
 - case 23
 - length 23
- narrowing conversions 35
- network access, java.net 123
- null values
 - case statements 85
 - in SQLJ functions 84
- nulls in Java-SQL 36–40
 - arguments to methods 38
 - using convert functions 39
- Number of Java Sockets, configuration parameter 130

O

- object mappable datatypes 98

Index

- obtaining connections 62
- options
 - external name 82
 - language java 82
 - parameter style java 82
- order by clauses 44
- ordering operations 44
- out parameter 90
- output mappable datatypes 98

P

- package names 145
- parameter style java parameter 88
- parameters
 - (Java VM) size of global fixed heap 125
 - (Java VM) size of process object heap 125
 - (Java VM) size of shared class heap 125
 - deterministic 88
 - external name 88
 - inout 90
 - input 90
 - language java 88
 - modifies sql data 88
 - not deterministic 88
 - output 90
 - parameter style java 88
- parentheses ()
 - in SQL statements xvii
- permissions
 - Java 6, 22
 - JDBC 59
 - SQLJ routines 77
- persistent data items 31
- PrintWriter, Java class 130
- procedure
 - creating SQLJ routine 76
 - enabling java.net 124
- procedures
 - client process, java.net 126
 - server process, java.net 126

Q

- questions and answers 4

R

- rearranging installed classes 20
- referencing
 - fields 27
- remove java command 19, 146
- removing classes 19
- removing JARs 19
- restrictions on Java in the database 9
- result sets 100
- ResultSet
 - mappable datatypes 98
- returns null on null input parameter, Java clause 82
- runtime
 - datatypes 36
- Runtime environment 11
- Runtime Java classes
 - location of 12
- runtime Java classes 12

S

- sample classes 51–53
 - address 51
 - address2Line 52
 - JDBCExamples 60–74
 - location of 10
 - misc 53
- saving text out of Adaptive server 125
- search order
 - function types 84
- security
 - SQLJ routines 77
- selecting Java objects 26
- serialization 146, 148
- server process 126
- server-side JDBC 6
- ServerSocket, Java class 124, 126
- set commands
 - allowed in Java methods 141
 - updating 43

- setting up 124
- shared class heap 124
- simply mappable datatypes 98
- Socket classes, using 125
- Socket, Java class 124
- sp_configure system procedure 13
- sp_depends system procedure 97
- sp_help system procedure 97
- sp_helpjava
 - syntax 18
 - utilitysp_helpjava 18
- sp_helpjava system procedure 97
- sp_helprotect system procedure 97
- SQL
 - expressions, include Java objects 6
 - function signature 81
 - procedure signature 87
 - wrappers 75, 79
- SQLJ create procedure command 87
- SQLJ functions 81–86
 - dropping 86
 - viewing information about 97
- SQLJ implementation
 - features not supported 103
 - features partially supported 103
 - SQLJ and Sybase differences 102
 - Sybase defined 103
- SQLJ standards 76
- SQLJ stored procedures 86–88, 96
 - capabilities of 86
 - deleting 96
 - modifying SQL data 88
 - using input and output parameters 90
 - viewing information about 97
- SQLJExamples class 105
- SQLJExamples.bestTwoEmps() method 78
- SQLJExamples.correctStates() method 78, 89
- SQLJExamples.job() method 78
- SQLJExamples.region() method 78, 83
- square brackets []
 - in SQL statements xvii
- standards for SQL 4
- standards specifications 4
- static methods 43, 78, 86
- static variables 46
- storage options

- in row 24
- String data
 - zero length 40
- string data 40
- style java keyword 88
- subtypes 34
- supertypes 34
- Sybase Central
 - creating a SQLJ function or procedure from 80
 - managing SQLJ procedures and functions from 80
 - viewing SQLJ routine properties from 81
- symbols
 - in SQL statements xvi, xvii
- syntax conventions
 - Java-SQL xvi
- syntax conventions, Transact-SQL xvi
- system procedures
 - helpjava 18
 - sp_depends 97
 - sp_help 97
 - sp_helpjava 97
 - sp_helprotect 97

T

- table definition 77
- temporary databases 50
- transact-SQL
 - commands, in Java methods 138
- transient data items 31

U

- unicode 40
- union operator 44
- updating Java objects 26
- URL
 - Java class 126
- URL class
 - accessing external server with XQL 128
 - downloading HTTP document 128
 - inserting data in a table 128
 - Java class 124, 128, 129
 - obtaining an HTTP document 128

Index

- sending email 128
- using 128
- URLConnection, Java class 124
- URLDecoder, Java class 124
- URLEncoder, Java class 124
- user-defined classes, creating 14
- using
 - Java and SQL together 6
 - Java classes 21, 50
 - Socket classes 125
 - URL class 126

V

- variable declarations 147
- variables 147
 - datatypes of 24
 - static 46
 - values assigned to 27
- viewing information
 - about installed classes 18
 - about installed JARs 18
- void methods 88

W

- where clause 35, 42, 45
- work databases 50

X

- XML
 - accessing documents with java.net 123

Z

- zero-length strings 40