



Java in Adaptive Server Enterprise

Adaptive Server Enterprise

12.5

DOCUMENT ID: 31652-01-1250-03

LAST REVISED: June 2001

Copyright © 1989-2001 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase database management 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, the Sybase logo, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Server IQ, Adaptive Warehouse, AnswerBase, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, Backup Server, ClearConnect, Client-Library, Client Services, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, E-Anywhere, E-Whatever, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, ImpactNow, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MySupport, Net-Gateway, Net-Library, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, RW-Library, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, Transact-SQL, Translation Toolkit, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server and XP Server are trademarks of Sybase, Inc. 3/01

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

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

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., 6475 Christie Avenue, Emeryville, CA 94608.

Contents

About This Book	xi
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	5
What are the key features?	5
How can I store Java instructions in the database?	6
How is Java executed in the database?	6
How can I use Java and SQL together?	7
What is the Java API?	8
How can I access the Java API from SQL?	8
Which Java classes are supported in the Java API?	8
Can I install my own Java classes?	9
Can I access data using Java?	9
Can I use the same classes on client and server?	9
How do I use Java classes in SQL?	9
Where can I find information about Java in the database?	10
What you cannot do with Java in the database	10
Sample Java classes	11
CHAPTER 2 Preparing for and Maintaining Java in the Database	13
The Java runtime environment	13
Java classes in the database	13
JDBC drivers	14
The Java VM	14
Configuring memory for Java in the database	15
Enabling the server for Java	15
Disabling the server for Java	15

Creating Java classes and JARs	16
Writing the Java code	16
Compiling Java code	16
Saving classes in a JAR file	16
Installing Java classes in the database	17
Using installjava	18
Referencing other Java-SQL classes	20
Viewing information about installed classes and JARs	20
Downloading installed classes and JARs	21
Removing classes and JARs	21
Retaining classes	21

CHAPTER 3	Using Java Classes in SQL.....	23
	General concepts	24
	Java considerations.....	24
	Java-SQL names.....	24
	Using Java classes as datatypes	25
	Creating and altering tables with Java-SQL columns.....	26
	Selecting, inserting, updating, and deleting Java objects.....	28
	Invoking Java methods in SQL	30
	Sample methods	31
	Exceptions in Java-SQL methods	31
	Representing Java instances	31
	Assignment properties of Java-SQL data items	32
	Datatype mapping between Java and SQL fields	35
	Character sets for data and identifiers	36
	Subtypes in Java-SQL data	36
	Widening conversions	36
	Narrowing conversions.....	37
	Runtime versus compile-time datatypes	38
	The treatment of nulls in Java-SQL data.....	38
	References to fields and methods of null instances	38
	Null values as arguments to Java-SQL methods	40
	Null values when using the SQL convert function	41
	Java-SQL string data	42
	Zero-length strings	42
	Type and void methods.....	43
	Java void instance methods	43
	Java void static methods	45
	Equality and ordering operations	45
	Evaluation order and Java method calls	46
	Columns	47
	Variables and parameters	47
	Static variables in Java-SQL classes	48

Java classes in multiple databases	49
Scope	49
Cross-database references	49
Inter-class transfers	50
Passing inter-class arguments	51
Temporary and work databases	51
Java classes	52

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	60
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 UpdateAction() method	63
Executing a select statement: the selectAction() method	64
Calling a SQL stored procedure: the callAction() method	65
Error handling in the native JDBC driver	67
The JDBCExamples class	69
The main() method	69
The internalMain() method	69
The connector() method	70
The doAction() method	71
The doSQL() method	72
The updateAction() method	72
The selectAction() method	73
The callAction() 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
SQLJExamples	78
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	87
Modifying SQL data.....	89
Using input and output parameters	90
Returning result sets	94
Viewing information about SQLJ functions and procedures	97
Advanced topics.....	98
Mapping Java and SQL datatypes	98
Using the command main method.....	102
SQLJ and Sybase implementation: a comparison	102
SQLJExamples class	105

CHAPTER 6	Introduction to XML in the Database	109
	Introduction	109
	Source code and javadoc.....	110
	References	110
	An overview of XML	111
	A sample XML document	111
	XML document types.....	116
	XSL: formatting XML information	118
	Character sets and XML data.....	118
	XML parsers	119

CHAPTER 7	Selecting Data with XQL	121
	Accessing the XML parser	121
	Setting the CLASSPATH environment variable	122
	Installing XQL in Adaptive Server	122
	Converting a raw XML document to a parsed version	123
	Inserting XML documents.....	123
	Updating XML documents	124
	Deleting XML documents	124
	Memory requirements for running the query engine inside Adaptive Server	124
	Using XQL	125
	Query structures that affect performance	127
	Examples.....	128
	Other usages of the XQL package.....	130
	com.sybase.xml.xql.XqlDriver syntax.....	130
	Validating your document.....	132
	Using XQL to develop standalone applications.....	133
	XQL methods	136
	Methods in com.sybase.xml.xql.Xql	136
	parse(String xmlDoc)	136

parse(InputStream xml_document, boolean validate)	137
query(String query, String xmlDoc)	138
query(String query, InputStream xmlDoc)	138
query(String query, SybXmlStream xmlDoc)	139
query(String query, JXml jxml)	139
sybase.aseutils.SybXmlStream	139
com.sybase.xml.xql.store.SybMemXmlStream	140
com.sybase.xml.xql.store.SybFileXmlStream	140
setParser(String parserName)	140
reSetParser	141

CHAPTER 8

Specialized XML Processing 143

The OrderXml class for order documents	143
OrderXml(String) constructor	144
OrderXml(date, customerid, server)	144
void order2Sql(String ordersTableName, String server)	144
static void createOrderTable (String ordersTableName, String server)	145
void setOrderElement (String elementName, String newValue)	145
String getItemElement (int itemNumber, String elementName)	145
void setItemElement (int itemNumber, String elementName, String newValue)	146
String getItemAttribute (int itemNumber, elementName, attributeName)	146
void setItemAttribute (int itemNumber, elementName, attributeName, newValue)	146
void appendItem (newItemid, newItemName, newQuantity, newUnit)	147
void deleteItem(int itemNumber)	147
Storing XML documents	147
Mapping and storage	148
Advantages and disadvantages of storage options	148
Client or server considerations	149
Creating and populating SQL tables for order data	149
Tables for element storage	150
Tables for document and hybrid storage	151
Using the element storage technique	151
Composing order documents from SQL data	151
Translating data from an XML order into SQL	153
Using the document storage technique	154
Storing XML order documents in SQL columns	154
Accessing the elements of stored XML order documents	155

	Server access to order elements.....	158
	Appending and deleting items in the XML document.....	159
	Using the hybrid storage technique.....	159
CHAPTER 9	XML for SQL Result Sets	161
	The ResultSetXML class.....	161
	ResultSetXml(String).....	161
	Constructor: ResultSetXml	
	(query, cdataColumns, colNames, server)	162
	ResultXml example	162
	String toSqlScript	
	(resultTableName, columnPrefix, goOption)	162
	String getColumn(int rowNumber, int columnNumber).....	163
	String getColumn(int rowNumber, String columnName)	163
	void setColumn	
	(int rowNumber, int columnNumber, newValue).....	163
	void setColumn	
	(int rowNumber, String columnName, newValue)	164
	Boolean allString	
	(int ColumnNumber, String compOp, String comparand)	164
	Boolean someString	
	(int columnNumber, String compOp, String comparand)	164
	A customizable example for different result sets.....	165
	The ResultSet document type	166
	Using the element storage technique.....	170
	Generating a ResultSet in the client.....	171
	Generating a result set in Adaptive Server	171
	Translating the XML ResultSet document in the client	172
	Translating the XML ResultSet Document in Adaptive Server.....	173
	Using the document storage technique.....	173
	Storing an XML ResultSet document in a SQL column.....	173
	Accessing the columns of stored ResultSet documents	174
	Quantified comparisons in stored ResultSet documents.....	177
CHAPTER 10	Debugging Java in the Database	181
	Introduction to debugging Java	181
	How the debugger works.....	181
	Requirements for using the Java debugger	181
	What you can do with the debugger.....	182
	Using the debugger.....	182
	Starting the debugger and connecting to the database.....	182
	Compiling classes for debugging	183
	Attaching to a Java VM	183

The Source window	184
Options	185
Setting breakpoints	186
Disconnecting from the database	188
A debugging tutorial	189
Before you begin	189
Start the Java debugger and connect to the database	189
Attach to a Java VM	190
Load source code into the debugger	190
Step through source code	191
Inspecting and modifying variables	192

CHAPTER 11

Network Access Using java.net	195
Overview	195
java.net classes	196
Setting up java.net	196
Example usage	197
Using socket classes	197
Using the URL class	200
User notes	202
Where to go for help	202

CHAPTER 12

Reference Topics	205
Assignments	205
Assignment rules at compile-time	206
Assignment rules at runtime	206
Allowed conversions	207
Transferring Java-SQL objects to clients	207
Supported Java API packages, classes, and methods	208
Supported Java packages and classes	208
Unsupported Java packages and classes	209
Unsupported java.sql methods and interfaces	209
Invoking SQL from Java	211
Special considerations	211
Transact-SQL commands from Java methods	212
Datatype mapping between Java and SQL	217
Java-SQL identifiers	219
Java-SQL class and package names	220
Java-SQL column declarations	221
Java-SQL variable declarations	221
Java-SQL column references	222
Java-SQL member references	223
Java-SQL method calls	224

Glossary	227
Index	233

About This Book

This book describes how to install and use Java classes and methods in the Sybase® Adaptive Server® Enterprise database.

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), the eXtensible Markup Language (XML), and the Extensible Query Language (XQL) 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, “Introduction to XML in the Database,” provides an overview of XML and the methods for storing XML documents in Adaptive Server and generating them from SQL data.

Chapters 7, 8, and 9 describe other ways that you can use XML in the Adaptive Server database.

- Chapter 7, “Selecting Data with XQL,” describes how to select raw data from Adaptive Server using the XQL language and display it as an XML document.

- Chapter 8, “Specialized XML Processing,” describes the OrderXML class, which is designed for an example application that uses XML documents for customer Order data, and is written specifically to process XML documents for order data.
- Chapter 9, “XML for SQL Result Sets,” describes the ResultSetXML class, which allows you to generate an XML document representing an SQL result set, and to access and update such an XML document.
- Chapter 10, “Debugging Java in the Database,” describes how you use the Sybase debugger with Java.
- Chapter 11, “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 12, “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.

Related documents

The following documents comprise the Sybase Adaptive Server Enterprise documentation:

- 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.
- *Configuring Adaptive Server Enterprise* for your platform – provides instructions for performing specific configuration tasks for Adaptive Server.
- *What's New in Adaptive Server Enterprise?* – describes the new features in Adaptive Server version 12.5, the system changes added to support those features, and the changes that may affect your existing applications.

- *Transact-SQL User's Guide* – documents Transact-SQL, Sybase's 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.
- *System Administration Guide* – provides in-depth information about administering servers and databases. This manual includes instructions and guidelines for managing physical resources, security, user and system databases, and specifying character conversion, international language, and sort order settings.
- *Reference Manual* – contains detailed information about all Transact-SQL commands, functions, procedures, and datatypes. This manual also contains a list of the Transact-SQL reserved words and definitions of system tables.
- *Performance and Tuning Guide* – explains how to tune Adaptive Server for maximum performance. This manual includes information about database design issues that affect performance, query optimization, how to tune Adaptive Server for very large databases, disk and cache issues, and the effects of locking and cursors on performance.
- *The Utility Guide* – documents the Adaptive Server utility programs, such as isql and bcp, which are executed at the operating system level.
- *The Quick Reference Guide* – provides a comprehensive listing of the names and syntax for commands, functions, system procedures, extended system procedures, datatypes, and utilities in a pocket-sized book. Available only in print version.
- *The System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Available only in print version.
- *Error Messages and Troubleshooting Guide* – explains how to resolve frequently occurring error messages and describes solutions to system problems frequently encountered by users.
- *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.
- *Using Sybase Failover in a High Availability System* – provides instructions for using Sybase's Failover to configure an Adaptive Server as a companion server in a high availability system.

- *Using Adaptive Server Distributed Transaction Management Features* – explains how to configure, use, and troubleshoot Adaptive Server DTM features in distributed transaction processing environments.
- *EJB Server User's Guide* – explains how to use EJB Server to deploy and execute Enterprise JavaBeans in Adaptive Server.
- *XA Interface Integration Guide for CICS, Encina, and TUXEDO* – provides instructions for using Sybase's DTM XA interface with X/Open XA transaction managers.
- *Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Sybase jConnect for JDBC Programmer's Reference* – describes the jConnect for JDBC product and explains how to use it to access data stored in relational database management systems.
- *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.
- *Historical Server User's Guide* – describes how to use Historical Server to obtain performance information for SQL Server and Adaptive Server.
- *Monitor Server User's Guide* – describes how to use Monitor Server to obtain performance statistics from SQL Server and Adaptive Server.
- *Monitor Client Library Programmer's Guide* – describes how to write Monitor Client Library applications that access Adaptive Server performance data.

Other sources of information

Use the Sybase Technical Library CD and the Technical Library Product Manuals Web site to learn more about your product:

- Technical Library CD contains product manuals and is included with your software. The DynaText browser (downloadable from Product Manuals at <http://www.sybase.com/detail/1,3693,1010661,00.html>) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

- Technical Library Product Manuals Web site is an HTML version of the Technical Library CD that you can access using a standard Web browser. In addition to product manuals, you will find links to the Technical Documents Web site (formerly known as Tech Info Library), the Solved Cases page, and Sybase/Powersoft newsgroups.

To access the Technical Library 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.

❖ For the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

❖ For the latest information on EBFs and Updates

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select EBFs/Updates. Enter user name and password information, if prompted (for existing Web accounts) or create a new account (a free service).
- 3 Specify a time frame and click Go.
- 4 Select a product.
- 5 Click an EBF/Update title to display the report.

❖ To create 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.

Java syntax conventions

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.jdbcx` 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 syntax conventions

This book uses the same font and syntax conventions for Transact-SQL as other Adaptive Server documents:

- Command names, command option names, utility names, utility flags, and other keywords are in Helvetica in paragraph text. For example:
`select` command
`isql` utility
`-f` flag
- Variables, or words that stand for values that you fill in, are in italics. For example:
user_name
server_name
- Code fragments are shown in a monospace font. Variables in code fragments (that is, words that stand for values that you fill in) are italicized. For example:

```
Connection con = DriverManager.getConnection  
("jdbc:sybase:Tds:host:port", props);
```


- You can disregard case when typing Transact-SQL keywords. For example, SELECT, Select, and select are the same.

Additional conventions for syntax statements in this manual are described in Table 1. Examples illustrating each convention can be found in the *System Administration Guide*.

Table 1: Syntax statement conventions

Key	Definition
{ }	Curly braces indicate that you choose at least one of the enclosed options. Do not include braces in your option.
[]	Brackets mean choosing one or more of the enclosed options is optional. Do not include brackets in your option.
()	Parentheses are to be typed as part of the command.
	The vertical bar means you may select only one of the options shown.
,	The comma means you may choose as many of the options shown as you like, separating your choices with commas to be typed as part of the command.

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	5
Sample Java classes	11

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 ANSI 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.

Adaptive Server allows you to:

- Generate XML-formatted documents from raw data stored in Adaptive Server.
- Store XML documents and data extracted from them in Adaptive Server.
- Query XML documents stored on the Web.

Adaptive Server uses the XML Query Language (XQL) to search XML documents. A Java-based XQL query processor is included with Adaptive Server. Because many of the tools commonly used to process XML are written in Java, Adaptive Server provides an excellent base for XML-SQL applications.

Standards

The SQLJ consortium of SQL vendors develops specifications for using Java with SQL. The consortium submits these specifications to ANSI for formal processing as standards. The standards can be found on the Web at <http://www.ansi.org>. In this document, SQLJ refers to capabilities compliant with SQLJ Part 1 of the standard specifications

Compliance with SQLJ standards ensures that Sybase functionality ports to all third-party, standards-compliant relational databases.

The standard specifications are in three parts:

- *Part 0* – “Database Language SQL—Part 10: Object Language Bindings (SQL/OLB),” ANSI X3.135.10-1998.

Specifications for embedding SQL statements in Java methods. Similar to the traditional SQL facilities for embedded SQL in COBOL and C and other languages. The Java classes containing embedded SQL statements are precompiled to pure Java classes with JDBC calls.

- *Part 1* – “SQLJ—Part 1: SQL Routines using the Java Programming Language,” ANSI NCITS N331.1.

Specifications for installing Java classes in a SQL system, and for invoking Java static methods as SQL stored procedures and functions.

- *Part 2* – “SQLJ—Part 2: SQL Types using the Java Programming Language,” ANSI NCITS N331.2.

Specifications for using Java classes as SQL datatypes.

Sybase supports Part 1 of the specification. In addition, Sybase extends the capabilities provided in the standard. For example, Adaptive Server allows you to reference Java methods and classes directly in SQL, whereas SQLJ Parts 1 and 2 require that you use SQL aliases.

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. See

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 12, “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 carry out these tasks 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 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/JavaSql` (UNIX) or `%SYBASE%\Ase-12_5\sample\JavaSql` (Windows NT). This subdirectory also contains Javadoc facilities so that you can view specifications about sample classes and methods with your Web browser.

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	13
Enabling the server for Java	15
Creating Java classes and JARs	16
Installing Java classes in the database	17
Viewing information about installed classes and JARs	20
Downloading installed classes and JARs	21
Removing classes and JARs	21

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 12, “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 version 1.2. It is compliant with and supports several classes and methods of JDBC version 2.0. See Chapter 12, “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 5.2 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 12, “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.

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.

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 21.

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 21 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 31.

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 `.` 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	24
Using Java classes as datatypes	25
Invoking Java methods in SQL	30
Representing Java instances	31
Assignment properties of Java-SQL data items	32
Datatype mapping between Java and SQL fields	35
Character sets for data and identifiers	36
Subtypes in Java-SQL data	36
The treatment of nulls in Java-SQL data	38
Java-SQL string data	42
Type and void methods	43
Equality and ordering operations	45
Evaluation order and Java method calls	46
Static variables in Java-SQL classes	48
Java classes in multiple databases	49
Java classes	52

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/JavaSql` (UNIX)
- `%SYBASE%\Ase-12_5\sample\JavaSql` (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 217.
- 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 219 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_addr 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')
```

n 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_add>>street, @zip = home_add>>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 43.

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_add>>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_addr>>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 217.

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_addr` 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 41 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 36, 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 38, 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(null, Address))>>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.

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 *J* is assigned to a target variable or column, or passed to a Java method, *J* 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/JavaSql`. (UNIX) or `%SYBASE%\Ase-12_5\sample\JavaSql` (Windows NT).

This is the Address class:

```
//
// Copyright (c) 1999
// Sybase, Inc
// Emeryville, CA 94608
// 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(street);
    }
}

```

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

```

//
// Copyright (c) 1999
// Sybase, Inc
// Emeryville, CA 94608
// 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) 1999
// Sybase, Inc
// Emeryville, CA 94608
// 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 "";
    }
}

```

Data Access Using JDBC

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	60
Error handling in the native JDBC driver	67
The JDBCExamples class	69

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 69 and in `$SYBASE/$SYBASE_ASE/sample/JavaSql` (UNIX) or `%SYBASE%\Ase-12_5\sample\JavaSql` (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 11. 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 *UpdateAction()* method

The `updateAction()` 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 (?). `updateAction()` 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, `updateAction()` executes the update statement.

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

Executing a *select* statement: the `selectAction()` method

The `selectAction()` 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, `selectAction()` supplies a value for the parameter marker after preparing the SQL statement:

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

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

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

Note For SQL statements that return no results, use `doSQL()` and `updateAction()`. 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.

`selectAction()` 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.

`selectAction()` 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 *callAction()* method

The `callAction()` 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*). `callAction()` 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.

`callAction()` 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 (?).

`callAction()` supplies values for the input parameters using JDBC `setInt()`, `setString()`, and `setObject()` methods that were used in the `doSQL()`, `updateAction()`, and `selectAction()` 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, `callAction()` 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);
```

`callAction()` then executes the call statement and obtains the output values using the same `getString()` and `getObject()` methods that the `selectAction()` 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 internalMain 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 internalMain() method

```
// A JDBCExamples method equivalent to 'main',
```

```
// to be called from SQL or Java in the server

public static String internalMain(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 += updateAction(con);
} else if (action.equals("select")) {
    workString += selectAction(con);
} else if (action.equals("call")) {
    workString += callAction(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 updateAction() method

// A method that updates a certain row of the 'xmp' table.

// This method illustrates prepared statements and parameter markers.

```
public static String updateAction(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 selectAction() 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 selectAction(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 callAction() 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 callAction(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
General issues	76
Using Sybase Central to manage SQLJ functions and procedures	80
SQLJ user-defined functions	81
SQLJ stored procedures	87
Viewing information about SQLJ functions and procedures	97
Advanced topics	98
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 104.

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.

- Only public static (class) methods can be referenced in a SQLJ function or stored procedure.
- Datatype mapping is checked when the SQLJ routine is created. During the execution of a SQLJ routine, data is passed from SQL to Java and back to SQL. Any data conversions required during execution must follow the rules for datatype mapping outlined in the JDBC standard.

Refer to “Mapping Java and SQL datatypes” on page 98 for a discussion of datatype mapping and conversions for SQLJ routines.

- If a method referenced by a SQLJ function or stored procedure invokes SQL, returns parameters from the SQL system to the calling environment, or returns result sets from SQL to the calling environment, you must use an Adaptive Server JDBC interface, such as Sybase jConnect or the internal Adaptive Server JDBC driver, that enables object-oriented access to the relational database.

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 30) 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*.

SQLJExamples

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)
```

he 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.

❖ **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.
- es 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 86.

- 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.

- Clauses `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.
- Clauses `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 98.

- 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_emps
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 98 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_emp 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.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());
    }
}
```

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_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;
}
```

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 98 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 99.

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		<code>java.lang.String</code>
nchar		<code>java.lang.String</code>
varchar/univarchar		<code>java.lang.String</code>
nvarchar		<code>java.lang.String</code>
text		<code>java.lang.String</code>
numeric		<code>java.math.BigDecimal</code>
decimal		<code>java.math.BigDecimal</code>
money		<code>java.math.BigDecimal</code>
smallmoney		<code>java.math.BigDecimal</code>
bit	<code>boolean</code>	<code>Boolean</code>
tinyint	<code>byte</code>	<code>Integer</code>
smallint	<code>short</code>	<code>Integer</code>
integer	<code>int</code>	<code>Integer</code>
real	<code>float</code>	<code>Float</code>
float	<code>double</code>	<code>Double</code>
double precision	<code>double</code>	<code>Double</code>
binary		<code>byte[]</code>
varbinary		<code>byte[]</code>
datetime		<code>java.sql.Timestamp</code>
smalldatetime		<code>java.sql.Timestamp</code>

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. When you execute the class from the command line as in the following:

```
java JDBCExamples
```

it is the command main method that executes.

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 SQL 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 sqlj.install_jar, sqlj.replace_jar, sqlj.remove_jar, and similar utilities to install, replace, and remove JAR files.	Supports the installjava utility and the remove java Transact-SQL 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.

SQLJ category	SQLJ standard	Sybase implementation
create procedure and create function commands	Supports the key words no sql, contains sql, reads sql data, and modifies sql data to specify the SQL operations the Java method can perform.	Supports modifies sql data only.
create procedure command	Supports java.sql.ResultSet and the SQL/OLB iterator declaration.	Supports java.sql.ResultSet only.
drop procedure and drop function commands	Supports the key word restrict, 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 restrict 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 deterministic not deterministic keywords, which specify whether or not the procedure or function always returns the same values for the out and inout parameters and the function result.	Supports only the syntax for deterministic not deterministic, not the functionality.
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 create 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 create command is executed, which enables faster execution.
create procedure and create function commands	Can specify the create procedure or create function commands within deployment descriptor files or as SQL DDL statements. The implementation defines which way (or ways) the commands are supported.	Supports create procedure and create function as SQL DDL statements outside of deployment descriptors.

SQLJ category	SQLJ standard	Sybase implementation
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.SQLException</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 drop procedure or drop function commands within deployment descriptor files or as SQL DDL statements. The implementation defines which way (or ways) the commands are supported.	Supports create procedure and create function 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/JavaSql`. (UNIX) or `%SYBASE%\Ase-12_5\sample\JavaSql` (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_emp 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());
        }
    }

    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;
    }
}
```

Introduction to XML in the Database

This chapter provides an overview of the eXtensible Markup Language (XML), and methods for storing XML documents in Adaptive Server and generating them from SQL data.

Topic	Page
Introduction	109
An overview of XML	111

Other XML topics are described in these chapters:

- Chapter 7, “Selecting Data with XQL,” describes how to select raw data from Adaptive Server using the XQL language and displaying it as an XML document.
- Chapter 8, “Specialized XML Processing,” describes the OrderXML class, which is designed for an example application that uses XML documents for customer order data, and is written specifically to process XML documents for order data.
- Chapter 9, “XML for SQL Result Sets,” describes the ResultSetXML class, which allows you to generate an XML document representing a SQL result set, and to access and update such an XML document.

Introduction

Like Hypertext Markup Language (HTML), XML is a markup language and a subset of Standardized General Markup Language (SGML). XML, however, is more complete and disciplined, and it allows you to define your own application-oriented markup tags. These properties make XML particularly suitable for data interchange.

You can generate XML-formatted documents from data stored in Adaptive Server and, conversely, store XML documents and data extracted from them in Adaptive Server. You can also use Adaptive Server to search XML documents stored on the Web.

Adaptive Server uses the XML Query Language (XQL) to search XML documents. XQL is a path-based query language that searches the XML documents using the XML structure.

Many of the XML tools needed to generate and process XML documents are written in Java. Java in Adaptive Server provides a good base for XML-SQL applications using both general and application-specific tools.

The XQL processor is a Java facility that is included with the Adaptive Server. It allows you to query and access XML data stored in Adaptive Server, and to display the result set as XML documents. See Chapter 7, “Selecting Data with XQL.”

Source code and javadoc

Adaptive Server includes the Java source code for the `XMLResultSet` and `OrderXML` classes. These classes provide an introduction for coding Java classes to process XML. The Java source code is in:

- `$SYBASE/ASE-12_5/sample/JavaSql` (UNIX)
- `%SYBASE%\Ase-12_5\sample\JavaSql` (Windows NT)

These directories also contain javadoc-generated HTML pages containing specifications for the referenced packages, classes, and methods.

References

This chapter presents an overview of XML. For detailed information, refer to these Web documents:

- World Wide Web Consortium (W3C), at <http://www.w3.org>
- W3C, Document Object Model (DOM), at <http://www.w3.org/DOM/>
- W3C, Extensible Markup Language (XML), at <http://www.w3.org/XML/>
- W3C, Extensible Stylesheet Language (XSL), at <http://www.w3.org/TR/WD-xsl/>

An overview of XML

XML is a markup language and subset of SGML. It was created to provide functionality that goes beyond that of HTML for Web publishing and distributed document processing.

XML is less complex than SGML, but more complex and flexible than HTML. Although XML and HTML can usually be read by the same browsers and processors, XML has characteristics that make it better able to share documents:

- XML documents possess a strict phrase structure that makes it easy to find and access data. For example, opening tags of all elements must have both an opening tag and a corresponding closing tag, for example, `<p> A paragraph.</p>`.
- XML lets you develop and use tags that distinguish different types of data, for example, customer numbers or item numbers.
- XML lets you create an application-specific document type, which makes it possible to distinguish one kind of document from another.
- XML documents allow different displays of the XML data. XML documents, like HTML documents, contain only markup and content; they do not contain formatting instructions. Formatting instructions are normally provided on the client using eXtensible Style Language (XSL) specifications.

You can store XML documents in Adaptive Server as:

- XML in a field of a Java object
- XML in a text or image column
- XML in a char or varchar column
- Parsed XML in an image column

A sample XML document

The sample Order document is designed for a purchase order application. Customers submit orders, which are identified by a date and a customer ID. Each order item has an item ID, an item name, a quantity, and a unit designation.

It might display on screen like this:

ORDER

Date: July 4, 1999

Customer ID: 123

Customer Name: Acme Alpha

Items:

Item ID	Item Name	Quantity
987	Coupler	5
654	Connector	3 dozen
579	Clasp	1

The following is one representation of this data in XML:

```
<?xml version="1.0"?>
  <Order>
    <Date>1999/07/04</Date>
    <CustomerId>123</CustomerId>
    <CustomerName>Acme Alpha</CustomerName>
    <Item>
      <ItemId> 987</ItemId>
      <ItemName>Coupler</ItemName>
      <Quantity>5</Quantity>
    </Item>
    <Item>
      <ItemId>654</ItemId>
      <ItemName>Connector</ItemName>
      <Quantity unit="12">3</Quantity>
    </Item>
    <Item>
      <ItemId>579</ItemId>
      <ItemName>Clasp</ItemName>
      <Quantity>1</Quantity>
    </Item>
  </Order>
```

The XML document has two unique characteristics:

- The XML document does not indicate type, style, or color for specifying item display.
- The markup tags are strictly nested. Each opening tag (*<tag>*) has a corresponding closing (*</tag>*).

The XML document for the order data consists of:

- The XML declaration, `<?xml version="1.0"?>`, identifying “Order” as an XML document.

XML represents documents as character data. In each document, you specify the character encoding (character set), either explicitly or implicitly. To explicitly specify the character set, include it in the XML declaration. For example:

```
<?xml version="1.0" encoding="ISO-8859-1">
```

If you do not include the character set in the XML declaration, the default, UTF8, is used.

Note When the default character sets of the client and server differ, Adaptive Server bypasses normal character-set translations so that the declared character set continues to match the actual character set. See “Character sets and XML data” on page 118.

- User-created element tags, such as `<Order>...</Order>`, `<CustomerId>...</CustomerId>`, `<Item>...</Item>`.
- Text data, such as “Acme Alpha,” “Coupler,” and “579.”
- Attributes embedded in element tags, such as `<Quantity unit = “12”>`. This embedding allows you to customize elements.

If your document contains these components, and the element tags are strictly nested, it is called a **well-formed XML document**. In the example above, element tags describe the data they contain, and the document contains no formatting instructions.

Here is another example of an XML document:

```
<?xml version="1.0"?>
<Info>
  <OneTag>1999/07/04</OneTag>
  <AnotherTag>123</AnotherTag>
  <LastTag>Acme Alpha</LastTag>
<Thing>
  <ThingId> 987</ThingId>
  <ThingName>Coupler</ThingName>
  <Amount>5</Amount>
  <Thing/>
<Thing>
  <ThingId>654</ThingId>
  <ThingName>Connector</ThingName>
</Thing>
```

```
<Thing>
  <ThingId>579</ThingId>
  <ThingName>Clasp</ThingName>
  <Amount>1</Amount>
</Thing>
</Info>
```

This example, called “Info,” is also a well-formed document and has the same structure and data as the XML Order document. However, it would not be recognized by a processor designed for Order documents because Info uses a different document type definition (DTD). For more information about DTDs, see “XML document types” on page 116).

HTML display of Order data

Consider a purchase order application. Customers submit orders, which are identified by a Date and the CustomerID, and which list one or more items, each of which has an ItemID, ItemName, Quantity, and units.

The data for such an order might be displayed on a screen as follows:

ORDER

Date: July 4, 1999

Customer ID: 123

Customer Name: Acme Alpha

Items:

Item ID	Item Name	Quantity
987	Coupler	5
654	Connector	3 dozen
579	Clasp	1

This data indicates that the customer named “Acme Alpha,” whose Customer Id is “123”, submitted an order on 1999/07/04 for couplers, connectors, and clasps.

The HTML text for this display of order data is as follows:

```
<html>
<body>
<p>ORDER
```

[illegible]

This HTML text has certain limitations:

- It contains both data and formatting specifications.
 - The data is the Customer Id, , and the various Customer Name, Item Names, and Quantities.
 - The formatting specifications are the indications for type style (*....*), color (*bcolor=white*), and layout (*<table>....</table>*), and also the supplementary field names, such as “*Customer Name*”, etc.
- The structure of HTML documents is not well suited for extracting data.

Some elements, such as tables, require strictly bracketed opening and closing tags, but other elements, such as paragraph tags (“*p*”), have optional closing tags.

Some elements, such as paragraph tags (“*p*”) are used for many sorts of data, so it is difficult to distinguish between a “123” that is a Customer ID and a “123” that is an Item ID, without specialized inference from surrounding field names.

This merging of data and formatting, and the lack of strict phrase structure, makes it difficult to adapt HTML documents to different presentation styles, and makes it difficult to use HTML documents for data interchange and storage. XML is similar to HTML, but includes restrictions and extensions that address these drawbacks.

XML document types

A **document type definition** (DTD) defines the structure of a class of XML documents, making it possible to distinguish between classes. A DTD is a list of element and attribute definitions unique to a class. Once you have set up a DTD, you can reference that DTD in another document, or embed it in the current XML document.

The DTD for XML Order documents, discussed in “A sample XML document” looks like this:

```
<!ELEMENT Order (Date, CustomerId, CustomerName,Item+)>
<!ELEMENT Date (#PCDATA)>
<!ELEMENT CustomerId (#PCDATA)>
<!ELEMENT CustomerName (#PCDATA)>
<!ELEMENT Item (ItemId, ItemName, Quantity)>
<!ELEMENT ItemId (#PCDATA)>
<!ELEMENT ItemName (#PCDATA)>
<!ELEMENT Quantity (#PCDATA)>
<!ATTLIST Quantity units CDATA #IMPLIED>
```

Read line by line, this DTD specifies that:

- An order must consist of a date, a customer ID, a customer name, and one or more items. The plus sign, “+”, indicates one or more items. Items signaled by a plus sign are required. A question mark in the same place indicates an optional element. An asterisk in the element indicates that an element can occur zero or more times. (For example, if the word “Item*” in the first line above were starred, there could be no items in the order, or any number of items.)
- Elements defined by “(#PCDATA)” are character text.
- The “<ATTLIST...>” definition in the last line specifies that quantity elements have a “units” attribute; “#IMPLIED”, at the end of the last line, indicates that the “units” attribute is optional.

The character text of XML documents is not constrained. For example, there is no way to specify that the text of a quantity element should be numeric, and thus the following display of data would be valid:

```
<Quantity unit="Baker's dozen">three</Quantity>
<Quantity unit="six packs">plenty</Quantity>
```

Restrictions on the text of elements must be handled by the applications that process XML data.

An XML's DTD must follow the `<?xml version="1.0"?>` instruction. You can either include the DTD within your XML document, or you can reference an external DTD.

- To reference a DTD externally, use something similar to:

```
<?xml version="1.0"?>
<!DOCTYPE Order SYSTEM "Order.dtd">
<Order>
...
</Order>
```

- Here's how an embedded DTD might look:

```
<?xml version="1.0"?>
<!DOCTYPE Order [
  <!ELEMENT Order (Date, CustomerId, CustomerName,
    Item+)>
  <!ELEMENT Date (#PCDATA)>
  <!ELEMENT CustomerId (#PCDATA)>
  <!ELEMENT CustomerName (#PCDATA)>
  <!ELEMENT Item (ItemId, ItemName, Quantity)>
  <!ELEMENT ItemId (#PCDATA)>
  <!ELEMENT ItemName (#PCDATA)>
  <!ELEMENT Quantity (#PCDATA)>
  <!--ATTLIST Quantity units CDATA #IMPLIED-->
]>
<Order>
  <Date>1999/07/04</Date>
  <CustomerId>123</CustomerId>
  <CustomerName>Acme Alpha</CustomerName>
  <Item>
    ...
  </Item>
</Order>
```

DTDs are not required for XML documents. However, a **valid XML document** has a DTD and conforms to that DTD.

XSL: formatting XML information

You can use XSL to format XML documents. XSL specifications (style sheets) consist of a set of rules that define the transformation of an XML document into either an HTML document or a different XML document:

- XSL specifications that transform an XML document into HTML can specify normal HTML formatting details in the output HTML.
- XSL specifications that transform an XML document into another XML document can map the input XML document to an output XML document with different element names and phrase structure.

You can create your own style sheets to display particular classes for particular applications. XSL is normally used with presentation applications rather than with applications for data interchange or storage.

Character sets and XML data

If the declared character sets of your client and server differ, you must be careful when declaring the character set of your XML documents.

Every XML document has a character-set value. If that encoding is not declared in the XML declaration, the default value of UTF8 is assumed. The XML processor, when parsing the XML data, reads this value and handles the data accordingly. When the default character set of the client and server differ, Adaptive Server bypasses normal character-set conversions to ensure that the declared character set and the actual character set remain the same.

- If you introduce an XML document into the database by providing the complete text in the values clause of an insert statement, Adaptive Server translates the entire SQL statement into the server's character set before processing the insertion. This is the way Adaptive Server normally translates character text, and you must make sure that the declared character set of the XML document matches that of the server.
- If you introduce an XML document into the database using `writetext` or Open Client CT-Library or Open Client DB-Library programs, Adaptive Server recognizes the XML document from the XML declaration and does *not* translate the character set to that of the server.
- If you read an XML document from the database, Adaptive Server does *not* translate the character set of the data to that of the client, since doing so might compromise the integrity of the XML document.

XML parsers

You can analyze XML documents and extract their data using SQL character-string operations, such as `substring`, `charindex`, and `patindex`. However, it is more efficient to use Java in SQL, and to use tools written in Java, such as XML parsers.

XML parsers can:

- Check that a document is well-formed and valid.
- Handle character-set issues.
- Generate a Java representation of a document's parse tree.
- Build or modify a document's parse tree.
- Generate a document's text from its parse tree.

Many XML parsers are available with a free license or are in the public domain. They normally implement two standard interfaces: the Simple API for XML (SAX) and the Document Object Model (DOM).

- *SAX* is an interface for parsing. It specifies input sources, character sets, and routines to handle external references. While parsing, it generates events so that user routines can process the document incrementally, and it returns a DOM object that is the parse tree of the document.
- *DOM* is an interface for the parse tree of an XML document. It provides facilities for stepping through and assembling a parse tree.

Applications that use the SAX and DOM interfaces are portable across XML parsers.

Selecting Data with XQL

This chapter describes how you use XQL to select raw data from Adaptive Server, using the XQL language and display the results as an XML document.

Topic	Page
Accessing the XML parser	121
Setting the CLASSPATH environment variable	122
Installing XQL in Adaptive Server	122
Other usages of the XQL package	130
XQL methods	136

Note isql displays only the first 50 characters of a result set that is derived from XML data. However, the examples in this chapter display the entire result set for purposes of illustration. To see the entire result set for any of the examples, use `com.sybase.xml.xql.XqlDriver` to run the query. See “Other usages of the XQL package” on page 130. You can also use the JDBC client, which helps you to store the result as a `java.lang.String`.

Adaptive Server features a query engine written in Java, which you can either install in the server, or run outside the server. Running it outside the server is like running any Java program on the command line.

This chapter first addresses running the query engine as a standalone program, outside Adaptive Server. See “Installing XQL in Adaptive Server” on page 122 for instructions on running the query engine inside Adaptive Server.

Accessing the XML parser

Whether you install your query engine as a standalone program or inside Adaptive Server, you must first access the XML parser. Sybase recommends the *xerces.jar* (vs.1.3.1) parser, available at

- `$SYBASE/ASE-12_5/lib/xerces.jar` (UNIX)
- `%SYBASE%\5\ASE-12_5\lib\xerces.jar` (Windows NT)

You can download the parser from:

Xerces Java Parser at <http://xml.apache.org/xerces-j>.

You can also use any parser that is compliant with SAX 2.0.

Setting the CLASSPATH environment variable

To create a standalone program outside Adaptive Server, you must set your CLASSPATH environment variable to include the directories that contain *xerces.jar* and *xml.zip*. For UNIX , enter:

```
setenv CLASSPATH $SYBASE/ASE-12_5/lib/xerces.jar
$SYBASE/ASE-12_5/lib/xml.zip
```

For Windows NT, enter:

```
set CLASSPATH = D:\%SYBASE%\ASE-12_5\lib\xerces.jar
D:\%SYBASE%\ASE-12_5\lib\xml.zip
```

Installing XQL in Adaptive Server

This section assumes you have already enabled Java in Adaptive Server. For information, see Chapter 2, “Preparing for and Maintaining Java in the Database.”

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

```
installjava
-f file_name
[-new | -update ]
...
```

Where:

- *file_name* is the name of the JAR file you are installing in the server.
- `new` informs the server this is a new file.

- update informs the server you are updating an existing JAR file.

For more information about `installjava`, see the Utility Guide.

To add support for XML in Adaptive Server, you must install the *xml.zip* and *xerces.jar* files. These files are located in the directories `$SYBASE/ASE-12_5/lib/xml.zip` and `$SYBASE/ASE-12_5/lib/xerces.jar`

For example, to install *xml.zip*, enter:

```
installjava -Usa -P -Sserver_name -f $SYBASE/ASE-12_5/lib/xml.zip
```

To install *xerces.jar*, enter:

```
installjava -Usa -P -Sserver_name -f $SYBASE/ASE-12_5/lib/xerces.jar
```

Note To install *xerces.jar* in a database, you must increase the size of tempdb by 10MB.

Converting a raw XML document to a parsed version

Use the `parse()` method to convert and parse a raw text or image XML document and store the result. Use the `alter table` command to convert the raw XML document. For example:

```
alter table XMLTEXT add xmldoc IMAGE null
update XMLTEXT
set xmldoc = com.sybase.xml.xql.Xql.query.parse(xmlcol)
```

This example converts the `xmlcol` column of the `XMLTEXT` table to parsed data and stores it in the `xmldoc` column.

Inserting XML documents

Use the `parse()` method to insert an XML document, which takes the XML document as the argument and returns `sybase.aseutils.SybXmlStream`.

Adaptive Server has an implicit mapping between image or text data and `InputStream`. You can pass image or text columns to `parse()` without doing any casting. The `parse()` UDF parses the document and returns `sybase.ase.SybXmlStream`, which Adaptive Server uses to write the data to the image column. Adaptive Server writes this data to image columns only, not to text columns. The following is an insert statement; where XMLDAT is an image column:

```
insert XMLDAT
values (..,
com.sybase.xml.xql.Xql.parse("<xmldoc></xmldoc>",..))
```

Updating XML documents

To update a document, delete the original data and then insert the new data. The number of updates to a document or portion of a document are infrequent compared to the number of reads. An update is similar to:

```
update XMLDAT
set xmldoc =
com.sybase.xml.xql.Xql.parse("<xmldoc></xmldoc>")
```

Deleting XML documents

Deleting an XML document is similar to deleting any text column. For example, to delete a table named XMLDAT, enter:

```
delete XMLDAT
```

Memory requirements for running the query engine inside Adaptive Server

Depending on the size of the XML data you want to select and present as an XQL document, you may need to increase memory. For a typical XML document of size 2K, Sybase recommends that you set the configuration parameters in Java Services to the values shown in Table 7-1. For more information on configuration parameters, see the *Sybase Adaptive Server System Administration Guide*.

Table 7-1: Java Services memory parameters

Section	Reset value
enable java	1
size of process object heap	5000
size of shared class heap	5000
size of global fixed heap	5000

Using XQL

XML Query Language (XQL) has been designed as a general-purpose query language for XML. XQL is a path-based query language for addressing and filtering the elements and text of XML documents, and is a natural extension to XSL syntax. XQL provides a concise, understandable notation for pointing to specific elements and for searching for nodes with particular characteristics. XQL navigation is through elements in the XML tree.

Note SQL and XQL are independent languages. The examples presented here apply to XQL only.

The most common XQL operators include:

- Child operator, / – indicates hierarchy. The following example returns `<book>` elements that are children of `<bookstore>` elements from the `xmlcol` column of the `xmlimage` table:

```
select
com.sybase.xml.xql.Xql.query("/bookstore/book",
xmlcol)
from xmlimage

<xql_result>
  <book style=autobiography>
    <title>S
```

- Descendant operator, // – indicates that the query searches through any number of intervening levels. That is, a search using the descendant operator finds an occurrence of an element at any level of the XML structure. The following query finds all the instances of `<emph>` elements that occur in an `<excerpt>` element:

```
select com.sybase.xml.xql.Xql.query
  ("/bookstore/book/excerpt//emph",xmlcol)
from xmlimage
```

```
<xql_result>
  <emph>I</emph>
</xql_result>
```

- Equals operator, = – specifies the content of an element or the value of an attribute. The following query finds all examples where “last-name = Bob”:

```
select com.sybase.xml.xql.Xql.query
  ("/bookstore/book/author[last-name='Bob']", xmlcol)
from xmlimage
```

```
<xql_result>
  <author>
    <first-name>Joe</first-name>
    <last-name>Bob</last-name>
    <award>Trenton Literary Review Honorable Mention</award></author>
  <author>
    <first-name>Mary</first-name>
    <last-name>Bob</last-name>
    <publication>Selected Short Stories of
    <first-name>Mary</first-name>
    <last-name>Bob</last-name></publication></author>
  <author>
    <first-name>Toni</first-name>
    <last-name>Bob</last-name>
    <degree from=Trenton U>B.A.</degree>
    <degree from=Harvard>Ph.D.</degree>
    <award>Pulizer</award>
    <publication>Still in Trenton</publication>
    <publication>Trenton Forever</publication></author>
</xql_result>
```

- Filter operator, [] – filters the set of nodes to its left, based on the conditions inside the brackets. This example finds any occurrences of authors whose first name is Mary that are listed in a book element:

```
select com.sybase.xml.xql.Xql.query
  ("/bookstore/book[author/first-name = 'Mary']", xmlcol)
from xmlimage
<xql_result>
  <book style=textbook>
    <title>History of Trenton</title>
    <author>
      <first-name>Mary</first-name>
      <last-name>Bob</last-name>
      <publication>Selected Short Stories of
```



```
<first-name>Mary</first-name>
<last-name>Bob</last-name></publication></author>
<price>55</price></book>
```

- Subscript operator, [*index_ordinal*] – finds a specific instance of an element. This example finds the second book listed in the XML document. Remember that XQL is zero-based, so it begins numbering at 0:

```
select com.sybase.xml.xql.Xql.query("/bookstore/book[1]", xmlcol)
from xmlimage
Query returned true and the result is
<xql_result>
  <book style=textbook>
    <title>History of Trenton</title>
    <author>
      <first-name>Mary</first-name>
      <last-name>Bob</last-name>
    <publication>Selected Short Stories of
      <first-name>Mary</first-name>
      <last-name>Bob</last-name></publication></author>
    <price>55</price></book>
</xql_result>
```

- Boolean expressions – you can use Boolean expressions within filter operators. For example, this query returns all <author> elements that contain at least one <degree> and one <award>:

```
select com.sybase.xml.xql.Xql.query
  ("/bookstore/book/author[degree and award]", xmlcol)
from xmlimage

<xql_result>
  <author>
    <first-name>Toni</first-name>
    <last-name>Bob</last-name>
    <degree from=Trenton U>B.A.</degree>
    <degree from=Harvard>Ph.D.</degree>
    <award>Pulizer</award>
    <publication>Still in Trenton</publication>
    <publication>Trenton Forever</publication></author>
</xql_result>
```

Query structures that affect performance

This section describes examples that use the XML query engine in different ways.

Examples

The placement of the where clause in a query affects processing. For example, this query selects all the books whose author's first name is Mary:

```
select com.sybase.xml.xql.Xql.query
      ("/bookstore/book[author/first-name = 'Mary']", xmlcol)
from XMLDAT
where
      com.sybase.xml.xql.Xql.query("/bookstore/book
      [author/first-name= 'Mary']", xmlcol) !=
      convert(com.sybase.xml.xql.Xql, null)>>EmptyResult
<xql_result ><book style="textbook">
  <title>History of Trenton</title>
  <author>
    <first-name>Mary</first-name>
    <last-name>Bob</last-name>
  </author>
  <publication>
    Selected Short Stories of
    <first-name>Mary</first-name>
    <last-name>Bob</last-name>
  </publication>
  <price>55</price>
</book></xql_result>
```

query() is invoked twice, once in the where clause and once in the select clause, which means the query executes twice and may be slow for large documents.

Thus, you can save the result set in an object while executing the query in the where clause and then restore the result in the select clause.

Or, you can write a class like HoldString, that concatenates the results obtained from every invocation of com.sybase.xml.xql.Xql.query(), for each XML document in each row:

```
declare @result HoldString
select @result = new HoldString()
select @result>>get()
from XMLDAT
where
      @result>>put (com.sybase.xml.xql.Xql.query
      ("/bookstore/book[author/first-name= 'Mary']",
      xmlcol)) !=
convert (com.sybase.xml.xql.Xql, null) >>EmptyResult
```

Sybase advises that you do not store the result set in the where clause. The query does not always execute the where clause, so trying to retrieve its result in the select clause may generate an erroneous result set. HoldString is an example class.

Because Adaptive Server stores each document in a column of a given row, when the query scans a set of rows in the where clause, more than one row may satisfy the search criteria. If this occurs, the query returns a separate XML result document for each qualified row. For example, if you create the following table:

```
create table XMLTAB ( xmlcol image)
insert XMLTAB values
  ( com.sybase.xml.xql.Xql.parse(<xml><A><B><C>c</C></B></A></xml>) );
insert XMLTAB values
  ( com.sybase.xml.xql.Xql.parse(<xml><D><E><C>c</C></E></D></xml>) );
```

Then execute this query:

```
select com.sybase.xml.xql.Xql.query("//C", xmlcol)
from XMLTAB
```

You would expect to get the following result set:

```
<xql_result>
<C>c</C>
<C>c</C>
</xql_result>
```

Instead, the result set returns the same row twice, once from the select clause and once from the where clause:

```
<xql_result>
<C>c</C>
</xql_result>

<xql_result>
<C>c</C>
</xql_result>
```

Other usages of the XQL package

Note Sybase does not support these usages of the XQL package. These usages require JDK 1.2 or higher.

You can query XML documents from the command line, using the standalone application `com.sybase.xml.xql.XqlDriver`.

You can use Java package methods provided in `com.sybase.xml.xql.Xql` to query XML documents in Java applications. You can also use these Java package methods to query XML documents in Adaptive Server 12.5, using the Java VM feature.

`com.sybase.xml.xql.XqlDriver` can parse and query only XML documents stored as files on your local system. You cannot use `com.sybase.xml.xql.XqlDriver` to parse or query XML documents stored in a database or over the network.

`com.sybase.xml.xql.XqlDriver` can be useful for developing XQL scripts and learning XQL. However, Sybase recommends that you use `com.sybase.xml.xql.XqlDriver` only as a standalone program, and not as part of another Java application, because `com.sybase.xml.xql.XqlDriver` includes a `main()` method. A Java program can only include one `main()` method, and if you include `com.sybase.xml.xql.XqlDriver` in another Java program that includes `main()`, the application attempts to implement both `main()` methods, which causes an error in Java.

Sybase recommends that applications use the `com.sybase.xml.xql.Xql` class to interface with the XML query engine. The methods of this class are specified in the section “Methods in `com.sybase.xml.xql.Xql`” on page 136.

`com.sybase.xml.xql.XqlDriver` syntax

The syntax for `com.sybase.xml.xql.XqlDriver` is:

```
java com.sybase.xml.xql.XqlDriver
-qstring XQL_query
-validate true | false
-infile string
-outfile string
-help
-saxparser string
```

Where:

- qstring specifies the XQL query you are running.
- validate checks the validity of the XML documents.
- infile is the XML document you are querying.
- outfile is the operating system file where you are storing the parsed XML document.
- help displays the com.sybase.xml.xql.XqlDriver syntax.
- saxparser specifies the name of a CLASSPATH parser that is compliant with SAX 2.0.

For information about XQL, see “Using XQL” on page 125.

Sample queries

This query selects all the book titles from *bookstore.xml*:

```
java com.sybase.xml.xql.XqlDriver -qstring "/bookstore/book/title"  
    -infile bookstore.xml
```

Query returned true and the result is

```
<xql_result>  
<title>Seven Years in Trenton</title>  
<title>History of Trenton</title>  
<title>Trenton Today, Trenton Tomorrow</title>  
</xql_result>
```

This example lists all the author’s first names from *bookstore.xml*. XQL uses a zero-based numbering system; that is, “0” specifies the first occurrence of an element in a file.

```
java com.sybase.xml.xql.XqlDriver  
    -qstring "/bookstore/book/author/first-name[0]"  
    -infile bookstore.xml
```

Query returned true and the result is

```
<xql_result>  
    <first-name>Joe</first-name>  
    <first-name>Mary</first-name>  
    <first-name>Toni</first-name>  
</xql_result>
```

The following example lists all the authors in *bookstore.xml* whose last name is “Bob”:

```
java com.sybase.xml.xql.XqlDriver
    -qstring "/bookstore/book/author[last-name='Bob']"
    -infile bookstore.xml
```

Query returned true and the result is

```
<xql_result>
  <author>
    <first-name>Joe</first-name>
    <last-name>Bob</last-name>
    <award>Trenton Literary Review Honorable Mention</award></author>
  <author>
    <first-name>Mary</first-name>
    <last-name>Bob</last-name>
    <publication>Selected Short Stories of
    <first-name>Mary</first-name>
    <last-name>Bob</last-name></publication></author>
  <author>
    <first-name>Toni</first-name>
    <last-name>Bob</last-name>
    <degree from=Trenton U>B.A.</degree>
    <degree from=Harvard>Ph.D.</degree>
    <award>Pulizer</award>
    <publication>Still in Trenton</publication>
    <publication>Trenton Forever</publication></author>
</xql_result>
```

Validating your document

The `valid` option invokes a parser that makes sure the XML document you are querying conforms to its DTD. Your standalone XML document must have a valid DTD before you run the `validate` option.

For example, this command makes sure the *bookstore.xml* document conforms to its DTD:

```
java com.sybase.xml.xql.XqlDriver -qstring "/bookstore" -validate
    -infile bookstore.xml
```

Using XQL to develop standalone applications

You can use XQL to develop standalone applications, JDBC clients, JavaBeans, and EJBs to process XML data. The `query()` and `parse()` methods in `com.sybase.xml.xql.Xql` enable you to query and parse XML documents. Because you can write standalone applications, you do not have to depend on Adaptive Server to supply the result set. Instead, you can query XML documents stored as operating system files or stored out on the Web.

Example standalone application

The following example uses the `FileInputStream()` query to read *bookstore.xml*, and the `URL()` method to read a Web page named *bookstore.xml* which contains information about all the books in the bookstore:

```
String result;
FileInputStream XmlFile = new FileInputStream("bookstore.xml");
if ((result =
        Xql.query("/bookstore/book/author/first-name", XmlFile))
    != Xql.EmptyResult )
{
    System.out.println(result);
}else{
    System.out.println("Query returned false\n");
}

URL _url = new URL("http://mybookstore/bookstore.xml");
if ((result =
        Xql.query("/bookstore/book/author/first-name",url.openStream()))
    != Xql.EmptyResult )
{
    System.out.println(result);
}else{
    System.out.println("Query returned false\n");
}
```

Example JDBC client

The following code fragment uses the `Xql.query` method to query the `xmlcol` column in the `XMLTEXT` file:

```
String selectQuery = "select xmlcol from XMLTEXT";
Statement stmt = _con.createStatement();
ResultSet rs = (SybResultSet)stmt.executeQuery(selectQuery);
```

```
String result;

InputStream is = null;
while ((rs != null) && (rs.next()))
{
    is = rs.getAsciiStream(1);
    result = Xql.query("/bookstore/book/author", is);
}
```

the following example assumes that the parsed XML data is stored in an image column of the XMLDOC table. Although this application fetches an image column as a binary stream, it does not parse this during the query because it identifies the content of this binary stream as a parsed XML document. Instead, the application creates a *SybXmlStream* instance from it and then executes the query. All this is done using the *Xql.query()* method, and does not require any input from the user.

```
String selectQuery = "select xmlcol from XMLDOC";
Statement stmt = _con.createStatement();
ResultSet rs = (SybResultSet)stmt.executeQuery(selectQuery);
InputStream is = null;
String result
while ((rs != null) && (rs.next()))
{
    is = rs.getBinaryStream(1);
    result = Xql.query("/bookstore/book/author/first-name", is));
}
```

Example EJB example

You can write EJB code fragments that serve as query engines on an EJB server.

The code fragment below includes an EJB called *XmlBean*. *XmlBean* includes the *query()* method, which allows you to query any XML document on the Web. In this component, *query()* first creates an *XmlDoc* object, then queries the document.

The remote interface looks like:

```
public interface XmlBean extends javax.ejb.EJBObject
{
    /**
     * XQL Method
```



```
    */
    public String XQL(String query, URL location) throws
java.rmi.RemoteException
;
}
```

The Bean implementation looks like:

```
public class XmlBean extends java.lang.Object implements
javax.ejb.SessionBean
{
    ....
    /**
     * XQL Method
     */
    public String XQL(String query, java.net.URL location) throws
java.rmi.RemoteException
{
    try {
        String result;

        if((result =
            Xql.query(query, location.openStream())) !=
            Xql.EmptyResult)
        {
            return (result);
        }else{
            return (null);
        }
    }catch(Exception e){
        throw new java.rmi.RemoteException(e.getMessage());
    }
}

....
}
```

And the client code looks like:

```
....
Context ctx = getInitialContext();
// make the instance of the class in Jaguar
XmlBeanHome _beanHome =
(XmlBeanHome)ctx.lookup("XmlBean");
XmlBean _xmlBean = (XmlBean)_beanHome.create();
URL u = new URL("http://mywebsite/bookstore.xml");
String res= xmlBean.XQL("/bookstore/book/author/first-name",u);
```

XQL methods

The XQL methods that Sybase supports and provides with Adaptive Server follow. For more information on these methods see the Web sites in the Reference section of Chapter 6, “Introduction to XML in the Database”.

- attribute
- comment
- element
- id
- node
- pi
- textNode
- textName
- text
- value

Methods in com.sybase.xml.xql.Xql

The following methods are specific to com.sybase.xml.xql.Xql.

parse(String xmlDoc)

Description	Takes a Java string as an argument and returns <i>SybXmlStream</i> . You can use this to query a document using XQL.
Syntax	<code>parse(String xml_document)</code> Where: <ul style="list-style-type: none">• String is a Java string.• <i>xml_document</i> is the XML document where the string is located.
Examples	The following example:

```
SybXmlStream xmlStream = Xql.parse("<xml>..</xml>");
```

Returns *SybXmlStream*.

Usage

The parser does not:

- Validate the document if a DTD is provided.
- Parse any external DTDs
- Perform any external links (for example, XLinks)
- Navigate through IDREFs

parse(InputStream xml_document, boolean validate)

Description

Takes an *InputStream* and a boolean flag as arguments. The flag indicates that the parser should validate the document according to a specified DTD. Returns *SybXmlStream*. You can use this to query a document using XQL.

Syntax

```
parse(InputStream xml_document, boolean validate)
```

Where:

- *InputStream* is an input stream.
- *xml_document* is the XML document where the input stream originates.

Examples

The following example

```
SybXmlStream is = Xql.parse(new  
FileInputStream("file.xml"), true);
```

Returns *SybXmlStream*.

Usage

- A true value in the flag indicates that the parser should validate the document according to the specified DTD.
- A false value in the flag indicates that the parser does not validate the document according to the specified DTD.
- The parser does not:
 - Parse any external DTDs
 - Perform any external links (for example, XLinks)
 - Navigate through IDREFs

query(String query, String xmlDoc)

Description Queries an XML document. Uses the XML document as the input argument.

Syntax query(**String query**, *String xmlDoc*)

Where:

- *String query* is the string you are searching for.
- *String xmlDoc* is the XML document you are querying.

Examples The following returns the result as a Java string:

```
String result= Xql.query("/bookstore/book/author",
    "<xml>...</xml>");
```

Usage Returns a Java string.

query(String query, InputStream xmlDoc)

Description Queries an XML document using an input stream as the second argument.

Syntax query(**String query**, *InputStream xmlDoc*)

Where:

- *String query* is the string you are searching for.
- *Input Stream xmlDoc* is the XML document you are querying.

Examples This example queries the bookstore for authors listed in *bookstore.Xql*.

```
FileInputStream xmlStream = new FileInputStream("doc.xml");
String result = Xql.query("/bookstore/book/author", xmlStream);
```

The following example queries an XML document on the Web using a URL as the search argument:

```
URL xmlURL = new URL("http://mywebsite/doc.xml");
String result = Xql.query("/bookstore/book/author", xmlURL.openStream());
```

Usage Returns a Java string.

query(String query, SybXmlStream xmlDoc)

Description Queries the XML document using a parsed XML document as the second argument.

Syntax `query(String query, SybXmlStream)`

Where:

- *String query* is the string you are searching for.
- *xmlDoc* is the parsed XML document you are querying.

Examples This example queries the bookstore for authors listed in *bookstore.Xml*.

```
SybXmlStream xmlStream = Xql.parse("<xml>..</xml>");  
String result = Xql.query("/bookstore/book/author", xmlStream);
```

query(String query, JXml jxml)

Description Queries an XML document stored in a JXML format.

Syntax `query(String query, JXml jxml)`

Where:

- *String query* is the string you are searching.
- *JXml jxml* is an object created from the classes located in `$$SYBASE/ASE-12_5/samples/`

Examples This example queries for authors in *bookstore.Xql*

```
JXml xDoc = new JXml("<xml>...</xml>");  
String result = Xql.query("/bookstore/book/author", xDoc);
```

Usage Allows you to execute a query on an JXML document using XQL.

sybase.aseutils.SybXmlStream

Description Defines an interface that an InputStream needs to access parsed XML data while querying.

Syntax `sybase.aseutils.SybXmlStream interface`

com.sybase.xml.xql.store.SybMemXmlStream

Description	Holds the parsed XML document in main memory, an implementation of SybXMLStream that Sybase provides.
Syntax	com.sybase.xml.xql.store.SybMemXmlStream
Usage	The parse() method returns an instance of SybMemXmlStream after parsing an XML document.

com.sybase.xml.xql.store.SybFileXmlStream

Description	Allows you to query a file in which you have stored a parsed XML document.
Syntax	com.sybase.xml.xql.store.SybFileXmlStream { <i>file_name</i> } Where <i>file_name</i> is the name of the file in which you stored the parsed XML document.
Examples	In the following, a member of the RandomAccessFile reads a file and positions the data stream:

```
SybXmlStream xis = Xql.parse("<xml>..</xml>");  
FileOutputStream ofs = new FileOutputStream("xml.data");  
( (SybMemXmlStream)xis).writeToFile(ofs);  
  
SybXmlStream is = new SybFileXmlStream("xml.data");  
String result = Xql.query("/bookstore/book/author", is);
```

setParser(String parserName)

Description	This static method specifies the parser that the parse method should use. You should make sure that the specified parser class is accessible through the CLASSPATH and is compliant with SAX 2.0.
Syntax	setParser (<i>String parserName</i>) Where <i>string</i> is the name of the parser class.
Examples	

```
Xql.setParser("com.yourcompany.parser")
```

reSetParser

Description This static method resets the parser to the default parser that Sybase supplies (*xerces.jar*, Version. 1.3.1).

Syntax reSetParser

Examples This example resets your parser to the Sybase default parser.

```
xql.resetParser()
```


When you store XML documents of a particular type in Adaptive Server, you may want to update them or to process them in specialized ways. One way is to write a Java class specifically designed for the updates and processing of that type of XML document. This chapter shows an example of such a Java class, designed for the OrderXML documents described in Chapter 6, “Introduction to XML in the Database.”

Topic	Page
The OrderXml class for order documents	143
Storing XML documents	147
Creating and populating SQL tables for order data	149
Using the element storage technique	151
Using the document storage technique	154
Using the hybrid storage technique	159

This section first describes the OrderXML class and its methods, and then provides a simple example that demonstrates how you can store XML documents or the data that they contain in an Adaptive Server database.

The source code and Javadoc for the OrderXML class is located in:

- `$SYBASE/ASE-12_5/sample/JavaSql` (UNIX)
- `%SYBASE%\ASE-12_5\sample\JavaSql` (Windows NT)

The *OrderXml* class for order documents

The examples in this section use the OrderXML class and its methods for basic operations on XML Order documents.

OrderXML is a subclass of the JXml class, which is specialized for XML Order documents. The OrderXML constructor validates the document for the Order DTD. Methods of the OrderXml class support referencing and updating the elements of the Order document.

OrderXml(String) constructor

Validates that the *String* argument contains a valid XML Order document, and then constructs an *OrderXml* object containing that document. For example, assume that “doc” is a Java string variable containing an XML Order document, perhaps one read from a file:

```
xml.order.OrderXml ox = new xml.order.OrderXml(doc);
```

OrderXml(date, customerid, server)

The parameters are all *String*.

This method assumes a set of SQL tables containing Order data. The method uses JDBC to execute a SQL query that retrieves Order data for the given *date* and *customerId*. The method then assembles an XML Order document with the data.

server identifies the Adaptive Server on which to execute the query.

- If you invoke the method in a client environment, specify the server name.
- If you invoke the method in Adaptive Server (in a SQL statement or in isql), specify either an empty string or the string “jdbc:default:connection”, which indicates that the query should be executed on the current Adaptive Server.

For example:

```
xml.order.OrderXml ox = new OrderXml("990704", "123",  
    "antibes:4000?user=sa");
```

void order2Sql(String ordersTableName, String server)

Extracts the elements of the Order document and stores them in a SQL table created by the *createOrderTable()* method. *ordersTableName* is the name of the target table. *server* is as described for the *OrderXml* constructor. For example, if *ox* is a Java variable of type *OrderXml*:

```
ox.order2Sql("current_orders", "antibes:4000?user=sa");
```

This call extracts the elements of the Order document contained in *ox*, and uses JDBC to insert the extracted elements into rows and columns of the table named *current_orders*.

```
static void createOrderTable(String ordersTableName, String server)
```

static void createOrderTable (String ordersTableName, String server)

Creates a SQL table with columns suitable for storing Order data: *customer_id*, *order_date*, *item_id*, *quantity*, and *unit*. *ordersTableName* is the name of the new table. The *server* parameter is as described for the *OrderXml* constructor. For example:

```
xml.order.OrderXml.createOrderTable
    ("current_orders", "antibes:4000?user=sa");
```

String getOrderElement(String elementName)

elementName is "Date," "CustomerId," or "CustomerName." The method returns the text of the element. For example, if *ox* is a Java variable of type *OrderXml*:

```
String customerId = ox.getOrderElement("CustomerId");
String customerName = ox.getOrderElement("CustomerName");
String date = ox.getOrderElement("Date");
```

void setOrderElement (String elementName, String newValue)

elementName is as described for *getOrderElement()*. The method sets that element to *newValue*. For example, if *ox* is a Java variable of type *OrderXml*:

```
ox.setOrderElement("CustomerName", "Acme Alpha Consolidated");
ox.setOrderElement("CustomerId", "987a");
ox.setOrderElement("Date", "1999/07/05");
```

String getItemElement (int itemNumber, String elementName)

itemNumber is the index of an item in the order. *elementName* is "ItemId," "ItemName," or "Quantity." The method returns the text of the item. For example, if *ox* is a Java variable of type *OrderXml*:

```
String itemId = ox.getItemElement(2, "ItemId");
String itemName = ox.getItemElement(2, "ItemName");
String quantity = ox.getItemElement(2, "Quantity");
```

void setItemElement (int itemNumber, String elementName, String newValue)

itemNumber and *elementName* are as described for the `getItemElement` method. `setItemElement` sets the element to *newValue*. For example, if *ox* is a Java variable of type `OrderXml`:

```
ox.setItemElement(2, "ItemId", "44");  
ox.setItemElement(2, "ItemName", "cord");  
ox.setItemElement(2, "Quantity", "3");
```

String getItemAttribute (int itemNumber, elementName, attributeName)

itemNumber and *elementName* are described as for `getItemElement()`. *elementName* and *attributeName* are both `String`. *attributeName* must be "unit." The method returns the text of the unit attribute of the item.

Note Since the Order documents currently have only one attribute, the *attributeName* parameter is unnecessary. It is included to illustrate the general case.

For example, if *ox* is a Java variable of type `OrderXml`

```
String itemid = ox.getItemAttribute(2, "unit")
```

void setItemAttribute (int itemNumber, elementName, attributeName, newValue)

itemNumber, *elementName*, and *attributeName* are as described for `getItemAttribute()`. *elementName*, *attributeName*, and *newValue* are `String`. The method sets the text of the unit attribute of the item to *newValue*. For example, if *ox* is a Java variable of type `OrderXml`:

```
ox.setItemAttribute(2, "unit", "13");
```

The parameters are all `String`. The method appends a new item to the document, with the given element values. For example, if *ox* is a Java variable of type `OrderXml`:

```
ox.appendItem("77", "spacer", "5", "12");
```

**void appendItem
(newItemId, newItemName, newQuantity, newUnit)**

The parameters are all String. The method appends a new item to the document, with the given element values. For example, if *ox* is a Java variable of type OrderXML:

```
ox.appendItem("77", "spacer", "5", "12");
```

void deleteItem(int itemNumber)

itemNumber is the index of an item in the order. The method deletes that item. For example, if *ox* is a Java variable of type OrderXml:

```
ox.deleteItem(2);
```

Storing XML documents

To use XML documents for data interchange in Adaptive Server, you must be able to store XML documents or the data that they contain in the database. To determine how best to accomplish this, consider the following:

- *Mapping and storage*: What sort of correspondence between XML documents and SQL data is most suitable for your system?
- *Client or server considerations*: Should the mapping take place on the client or the server?
- *Accessing XML in SQL*: How do you want to access the elements of an XML document in SQL?

The rest of this section discusses each of these considerations; the remainder of the chapter discusses the classes and methods you can use with XML, including:

- A simple example to illustrate the basics of data storage and exchange of XML documents
- A generalized example that you can customize for your own XML documents

Mapping and storage

There are three basic ways to store XML data in Adaptive Server: **element storage**, **document storage**, or **hybrid storage**, which is a mixture of both.

- Element storage – in this method, you extract data elements from an XML document and store them as data rows and columns in Adaptive Server.

For example, using the XML Order document, you can create SQL tables with columns for the individual elements of an order: Date, CustomerId, CustomerName, ItemId, ItemName, Quantity, and Units. You can then manage that data in SQL with normal SQL operations:

- To produce an XML document for Order data contained in SQL, retrieve the data, and assemble an XML document with it.
- To store an XML document with new Order data, extract the elements of that document, and update the SQL tables with that data.
- Document storage – in this method, you store an entire XML document in a single SQL column.
- For example, using the Order document, you can create one or more SQL tables having a column for Order documents. The datatype of that column could be:
 - SQL text, or
 - A generic Java class designed for XML documents, or
 - Java class designed specifically for XML Order documents
- Hybrid storage – in this method, you store an XML document in a SQL column, and also extract some of its data elements into separate columns for faster and more convenient access.

Again, using the Order example, you can create SQL tables as you would for document storage, and then include (or later add) one or more columns to store elements extracted from the Order documents.

Advantages and disadvantages of storage options

Each storage option has advantages and disadvantages. You must choose the option or options best for your operation.

- If you use element storage, all of the data from the XML document is available as normal SQL data that you can query and update using SQL operations. However, element storage has the overhead of assembling and disassembling the XML documents for interchange.
- Document storage eliminates the need for assembling and disassembling the data for interchange. However, you need to use Java methods to reference or update the elements of the XML documents while they are in SQL, which is slower and less convenient than the direct SQL access of element storage.
- Hybrid storage balances the advantages of element storage and document storage, but has the cost and complexity of redundant storage of the extracted data.

Client or server considerations

You can execute Java methods either on a client or on a server, which is a consideration for element storage and hybrid storage. Document storage involves little or no processing of the document.

- Element storage – if you map individual elements of an XML document to SQL data, in most cases the XML document is larger than the SQL data. It is generally more efficient to assemble and disassemble the XML document on the client and transfer only the SQL data between the client and the server.
- Hybrid storage – if you store both the complete XML document and extracted elements, then it is generally more efficient to extract the data from the server, rather than transfer it from the client.

Creating and populating SQL tables for order data

In this section, we create several tables that are designed to contain data from XML Order documents, so that we can demonstrate techniques for element, document, and hybrid data storage.

Tables for element storage

The following statements create SQL tables customers, orders, and items, whose columns correspond with the elements of the XML Order documents.

```
create table customers
  (customer_id varchar(5) not null unique,
   customer_name varchar(50) not null)

create table orders
  (customer_id varchar(5) not null,
   order_date datetime not null,
   item_id varchar(5) not null,
   quantity int not null,
   unit smallint default 1)

create table items
  (item_id varchar(5) unique,
   item_name varchar(20))
```

These tables need not have been specifically created to accommodate XML Order documents.

The following SQL statements populate the tables with the data in the example XML Order document (see “A sample XML document” on page 111):

```
insert into customers values("123", "Acme Alpha")

insert into orders values ("123", "1999/05/07",
  "987", 5, 1)

insert into orders values ("123", "1999/05/07",
  "654", 3, 12)

insert into orders values ("123", "1999/05/07",
  "579", 1, 1)

insert into items values ("987", "Widget")

insert into items values ("654", "Medium connector")

insert into items values ("579", "Type 3 clasp")
```

Use select to retrieve data from the tables:

```
select order_date as Date, c.customer_id as CustomerId,
  customer_name as CustomerName,
  o.item_id as ItemId, i.item_name as ItemName,
  quantity as Quantity, o.unit as unit
from customers c, orders o, items i
where c.customer_id=o.customer_id and
  o.item_id=i.item_id
```


Date	CustomerId	CustomerName	ItemId	ItemName	Quantity	Unit
July 4 1999	123	Acme Alpha	987	Coupler	5	1
July 4 1999	123	Acme Alpha	654	Connector	3	12
July 4 1999	123	Acme Alpha	579	Clasp	1	1

Tables for document and hybrid storage

The following SQL statement creates a SQL table for storing complete XML Order documents, either with or without extracted elements (for hybrid storage).

```
create table order_docs
  (id char(10) unique,
   customer_id varchar(5) null,
   -- For an extracted "CustomerId" element
   order_doc xml.order.OrderXml)
```

Using the element storage technique

This section describes the element storage technique for bridging XML and SQL.

- “Composing order documents from SQL data” on page 151 discusses the composition of an XML Order document from SQL data.
- “Translating data from an XML order into SQL” on page 153 discusses the decomposition of an XML Order document to SQL data.

Composing order documents from SQL data

In this example, Java methods generate an XML Order document from the SQL data in the tables created in “Creating and populating SQL tables for order data” on page 149.

A constructor method of the OrderXml class maps the data. An example call of that constructor might be:

```
new xml.order.OrderXml("990704", "123",
  "antibes:4000?user=sa");
```

This constructor method uses internal JDBC operations to:

- Execute a SQL query for the order data
- Generate an XML Order document with the data
- Return the OrderXml object that contains the Order document

You can invoke the OrderXml constructor in the client or the Adaptive Server.

- If you invoke the OrderXml constructor in the client, the JDBC operations that it performs use jConnect to connect to the Adaptive Server and perform the SQL query. It then reads the result set of that query and generates the order document on the client.
- If you invoke the OrderXml constructor in Adaptive Server, the JDBC operations that it performs use the native JDBC driver to connect to the current Adaptive Server and perform the SQL query. It then reads the result set and generates the Order document in Adaptive Server.

Generating an order on the client

Designed to be implemented on the client, main() invokes the constructor of the OrderXML class to generate an XML Order from the SQL data. That constructor executes a select for the given date and customer ID, and assembles an XML Order document from the result.

```
import java.io.*;
import util.*;
public class Sql2OrderClient {
    public static void main (String args[]) {
        try{
            xml.order.Order order =
                new xml.order.OrderXml("990704", "123","antibes:4000?user=sa");
            FileUtil.string2File("Order-sql2Order.xml",order.getXmlText());
        } catch (Exception e) {
            System.out.println("Exception:");
            e.printStackTrace();
        }
    }
}
```

Generating an order on the server

Designed for the server environment, the following SQL script invokes the constructor of the OrderXml class to generate an XML order from the SQL data:

```
declare @order xml.order.OrderXml
select @order =
    new xml.order.OrderXml('990704', '123','')
insert into order_docs (id, order_doc) values("3",
@order)
```

Translating data from an XML order into SQL

In this section, you extract elements from an XML Order document and store them in the rows and columns of the Orders tables. The examples illustrate this procedure in both server and client environments.

You translate the elements using the Java method `order2Sql()` of the `OrderXml` class. Assume that `xmlOrder` is a Java variable of type `OrderXml`:
`xmlOrder.order2Sql("orders_received", "antibes:4000?user=sa");`

The `order2Sql()` call extracts the elements of the XML Order document contained in variable `xmlOrder`, and then uses JDBC operations to insert that data into the `orders_received` table. You can call this method on the client or on Adaptive Server:

- Invoked from the client, `order2Sql()` extracts the elements of the XML Order document in the client, uses `jConnect` to connect to the Adaptive Server, and then uses the Transact-SQL insert command to place the extracted data into the table.
- Invoked from the server, `order2Sql()` extracts the elements of the XML Order document in the Adaptive Server, uses the native JDBC driver to connect to the current Adaptive Server, and then use the Transact-SQL insert command to place the extracted data into the table.

Translating the XML document on the client

Invoked from the client, the `main()` method of the `Order2SqlClient` class creates a table named `orders_received` with columns suitable for order data. It then extracts the elements of the XML order contained in the file *Order.xml* into rows and columns of *orders received*. It performs these actions with calls to static method `OrderXml.createOrderTable()` and instance method `order2Sql()`.

```
import util.*;
import xml.order.*;
import java.io.*;
import java.sql.*;
import java.util.*;
```

```
public class Order2SqlClient {
    public static void main (String args[]) {
        try{
            String xmlOrder =
                FileUtil.file2String("order.xml");
            OrderXml.createOrderTable("orders_received",
                "antibes:4000?user=sa");
            xmlOrder.order2Sql("orders_received",
                "antibes:4000?user=sa");
        } catch (Exception e) {
            System.out.println("Exception:");
            e.printStackTrace();
        }
    }
}
```

Translating the XML document on the server

Invoked from the server, the following SQL script invokes the OrderXml constructor to generate an XML Order document from the SQL tables, and then invokes the method OX.sql2Order(), which extracts the Order data from the generated XML and inserts it into the orders_received table.

```
declare @xmlorder OrderXml
select @xmlorder = new OrderXml('19990704', '123','')
select  @xmlorder>>order2Sql('orders_received', '')
```

Using the document storage technique

When using the document storage technique, you store a complete XML document in a single SQL column. This approach avoids the cost of mapping the data between SQL and XML when documents are stored and retrieved, but access to the stored elements can be slow and inconvenient.

Storing XML order documents in SQL columns

This section provides examples of document storage from the client and from the server.

Inserting an order document from a client file

The following command-line call is representative of how you can insert XML data into Adaptive Server from a client file. It copies the contents of the *Order.xml* file (using the `-I` parameter) to the Adaptive Server and executes the SQL script (using the `-Q` parameter) using the contents of *Order.xml* as the value of the question-mark (?) parameter.

```
java util.FileUtil -A putstring -I "Order.xml" \  
-Q "insert into order_docs (id, order_doc) \  
    values ('1', new xml.order.OrderXml(?)) " \  
-S "antibes:4000?user=sa"
```

Note The constructor invocation `new xml.order.OrderXml` validates the XML Order document.

Inserting a generated order document on the server

Executed on the server, the following SQL command generates an XML Order document from SQL data, and immediately inserts the generated XML document into the column of the `order_docs` table.

```
insert into order_docs (ID, order_doc)  
    select "2", new xml.order.OrderXml("990704", "123", "")
```

Accessing the elements of stored XML order documents

We have created a table named `order_docs` with a column named `order_doc`. The datatype of the `order_doc` column is `OrderXml`, a Java class that contains an XML Order document.

The `OrderXml` class contains several instance methods that let you reference and update elements of the XML Order document. They are described in “The `OrderXml` class for order documents” on page 143.

This section uses these methods to update the order document.

```
<?xml version="1.0"?>  
<!DOCTYPE Order SYSTEM "Order.dtd">  
<Order>  
    <Date>1999/07/04</Date>  
    <CustomerId>123</CustomerId>  
    <CustomerName>Acme Alpha</CustomerName>  
    <Item>
```

```
        <ItemId> 987</ItemId>
        <ItemName>Coupler</ItemName>
        <Quantity>5</Quantity>
    </Item><Item>
        <ItemId>654</ItemId>
        <ItemName>Connector</ItemName>
        <Quantity unit="12">3</Quantity>
    </Item><Item>
        <ItemId>579</ItemId>
        <ItemName>Clasp</ItemName>
        <Quantity>1</Quantity>
    </Item>
</Order>
```

Each XML Order document has exactly one *Date*, *CustomerId*, and *CustomerName*, and zero or more *Items*, each of which has an *ItemId*, *ItemName*, and *Quantity*.

Client access to order elements

The `main()` method of the `OrderElements` class is executed on the client. It reads the *Order.xml* file into a local variable, and constructs an `OrderXml` document from it. The method then extracts the “header” elements (*Date*, *CustomerId*, and *CustomerName*) and the elements of the first item of the order, prints those elements, and finally updates those elements of the order with new values.

```
import java.io.*;
import util.*;
public class OrderElements {
    public static void main ( String[] args) {
        try{

            String xml = FileUtil.file2String("Order.xml");
            xml.order.OrderXml ox =
                new xml.order.OrderXml(xml);

            // Get the header elements
            String cname = ox.getOrderElement("CustomerName");
            String cid   = ox.getOrderElement("CustomerId");
            String date  = ox.getOrderElement("Date");

            // Get the elements for item 1 (numbering from 0)
            String iName1 = ox.getItemElement(1, "ItemName");
            String iId1   = ox.getItemElement(1, "ItemId");
            String iQ1    = ox.getItemElement(1, "Quantity");
            String iU     = ox.getItemAttribute(1, "Quantity", "unit");
            System.out.println("\nBEFORE UPDATE: ")
```

```

System.out.println("\n    "+date+ "    "+ cname + "    "+cid);
System.out.println("\n    "+ iName1+" "+iId1+" "
    + iQ1 + "    " + iU + "\n");

// Set the header elements
ox.setOrderElement("CustomerName", "Best Bakery"
ox.setOrderElement("CustomerId", "531");
ox.setOrderElement("Date", "1999/07/31");

// Set the elements for item 1 (numbering from 0)
ox.setItemElement(1, "ItemName", "Flange");
ox.setItemElement(1, "ItemId", "777");
ox.setItemElement(1, "Quantity", "3");
ox.setItemAttribute(1, "Quantity", "unit", "13");

//Get the updated header elements
cname = ox.getOrderElement("CustomerName");
cid   = ox.getOrderElement("CustomerId");
date  = ox.getOrderElement("Date");

// Get the updated elements for item 1
// (numbering from 0)
iName1 = ox.getItemElement(1, "ItemName");
iId1   = ox.getItemElement(1, "ItemId");
iQ1    = ox.getItemElement(1, "Quantity");
iU     = ox.getItemAttribute(1, "Quantity", "unit");

System.out.println("\nAFTER UPDATE: ");
System.out.println("\n    "+date+ "    "+ cname + "    "+cid);
System.out.println("\n    "+ iName1+" "+iId1+" "
    + iQ1 + "    " + iU + "\n");

//Copy the updated document to another file
FileUtil.string2File("Order-updated.xml", ox.getXmlText())

} catch (Exception e) {
System.out.println("Exception:");
e.printStackTrace();
}
}
}

```

After implementing the methods in `OrderElements`, the order document stored in *Order-updated.xml* is:

```

<?xml version="1.0"?>
<!DOCTYPE Order SYSTEM 'Order.dtd'>
<Order>
  <Date>1999/07/31</Date>
  <CustomerId>531</CustomerId>

```

```
<CustomerName>Best Bakery</CustomerName>
<Item>
  <ItemId> 987</ItemId>
  <ItemName>Coupler</ItemName>
  <Quantity>5</Quantity>
</Item>
<Item>
  <ItemId>777</ItemId>
  <ItemName>Flange</ItemName>
  <Quantity unit="13">3</Quantity>
</Item>
<Item>
  <ItemId>579</ItemId>
  <ItemName>Clasp</ItemName>
  <Quantity>1</Quantity>
</Item>
</Order>
```

Server access to order elements

The preceding example showed uses of get and set methods in a client environment. You can also call those methods in SQL statements in the server:

```
select order_doc>>getOrderElement("CustomerId"),
       order_doc>>getOrderElement("CustomerName"),
       order_doc>>getOrderElement("Date")
from order_docs

select order_doc>>getItemElement(1, "ItemId"),
       order_doc>>getItemElement(1, "ItemName"),
       order_doc>>getItemElement(1, "Quantity"),
       order_doc>>getItemAttribute(1, "Quantity", "unit")
from order_docs

update order_docs
set order_doc = order_doc>>setItemElement(1, "ItemName",
      "Wrench")

update order_docs
set order_doc = order_doc>>setItemElement(2, "ItemId", "967")

select order_doc>>getItemElement(1, "ItemName"),
       order_doc>>getItemElement(2, "ItemId")
from order_docs

update order_docs
set order_doc = order_doc>>setItemAttribute(2, "Quantity",
```



```

        "unit", "6")
select order_doc>>getItemAttribute(2, "Quantity", "unit")
from order_docs

```

Appending and deleting items in the XML document

The Order class provides methods for adding and removing items from the Order document.

You can append a new item to the Order document with the `appendItem()` method, whose parameters specify `ItemId`, `ItemName`, `Quantity`, and units for the new item:

```

update order_docs
set order_doc = order_doc>>appendItem("864",
    "Bracket", "3", "12")

```

`appendItem()` is a void method that modifies the instance. When you invoke such a method in an update statement, you reference it as shown, as if it were an order-valued method that returns the updated item.

Delete an existing item from the order document using `deleteItem()`. `deleteItem()` specifies the number of the item to be deleted. The numbering begins with zero, so the following command deletes the second item from the specified row.

```

update order_docs
set order_doc = order_doc>>deleteItem(1)
where id = "1"

```

Using the hybrid storage technique

In the hybrid storage technique, you store the complete XML document in a SQL column and, at the same time, store elements of that document in separate columns. This technique often balances the advantages and disadvantages of element and document storage.

“Using the document storage technique” on page 154 demonstrates how to store the entire XML order document in the single column `order_docs.order_doc`. Using document storage, you must reference and access the *CustomerId* element in this way:

```
select order_doc>>getOrderElement("CustomerId") from order_docs
where order_doc>>getOrderElement("CustomerId") > "222"
```

To access *CustomerId* more quickly and conveniently than with the method call, but without first translating the Order into SQL rows and columns:

- 1 Add a column to the `order_docs` table for the *customer_id*:

```
alter table order_docs
add customer_id varchar(5) null
```

- 2 Update that new column with extracted *customerId* values.

```
update order_docs
set customer_id =
order_doc>>getOrderElement("CustomerId")
```

- 3 Reference *CustomerId* values directly:

```
select customer_id from order_docs
where customer_id > "222"
```

You can also define an index on the column.

Note This technique does not synchronize the extracted *customer_id* column with the *CustomerId* element of the *order_doc* column if you update either value.

This chapter describes the `ResultSetXml` class, which allows you to generate an XML document representing an SQL result set, and to access and update such an XML document.

Topic	Page
The <code>ResultSetXML</code> class	161
Generating a <code>ResultSet</code> in the client	171
Generating a result set in Adaptive Server	171

The source code for the `ResultSetXml` class is in the following directories:

- `$SBASE/ASE-12_5/sample/JavaSql` (UNIX)
- `%SYBASE%\ASE-12_5\sample\JavaSql` (Windows NT)

You can use the `ResultSetXML` class to process SQL result sets with XML and as an example of how to write Java code for accessing XML. Chapter 8, “Specialized XML Processing,” provides an additional example of Java code.

The `ResultSetXML` class

The `ResultSetXml` class is a subclass of the `JXml` class, which validates a document with the XML `ResultSet` DTD, and also provides methods for accessing and updating the elements of the contained XML `ResultSet` document.

`ResultSetXml(String)`

Validates that the argument contains a valid XML `ResultSet` document and constructs a *ResultSetXml* object containing that document. For example, if *doc* is a Java `String` variable containing an XML `ResultSet` document, read from a file:

```
xml.resultset.ResultSetXml rsx =  
    new xml.resultset.ResultSetXml (doc) ;
```

Constructor: **ResultSetXml** (**query**, **cdataColumns**, **colNames**, **server**)

- The parameters are all String.
- *query* is any SQL query that returns a result set.
- *server* identifies the Adaptive Server on which to execute the query. If you invoke the method in a client environment, specify the server name.

If you invoke the method in an Adaptive Server (in a SQL statement or isql), specify either an empty string or the string “jdbc:default:connection,” indicating that the query should be executed on the current Adaptive Server

- *cdata columns* indicates which columns should be XML CDATA sections.
- *colNames* indicates whether the resulting XML should specify “name” attributes in the “Column” elements

ResultSetXml example

The method connects to the server, executes the query, retrieves the SQL result set, and constructs a *ResultSetXml* object with that result set.

For example:

```
xml.resultset.ResultSetXml rsx =  
    new xml.resultset.ResultSetXml  
    ("select 1 as 'a', 2 as 'b', 3 ", "none", "yes",  
    "antibes:4000?user=sa");
```

This constructor call connects to the server specified in the last argument, evaluates the SQL query given in the first argument, and returns an XML *ResultSet* containing the data from the result set of the query.

String toSqlScript (**resultTableName**, **columnPrefix**, **goOption**)

- The parameters are all String.

- *resultTableName* is the table name for the create and insert statements. (SQL result sets do not specify a table name because they may be derived from joins or unions.)
- *columnPrefix* is the prefix to use in generated column names, which are needed for unnamed columns in the result set.
- *goOption* indicates whether the script is to include the go commands, which are required in isql, not in JDBC.

The method returns a SQL script with a create statement and a list of insert statements that create the result set data again.

For example, if *rsx* is a Java variable of type `ResultSetXml`:

```
rsx>>toSqlScript("systypes_copy", "column_", "yes")
```

String getColumn(int rowNumber, int columnNumber)

- *rowNumber* is the index of a row in the result set.
- *columnNumber* is the index of a column of the result set. The method returns the text of the specified column.

For example, if *rsx* is a Java variable of type `ResultSetXml`:

```
select rsx>>getColumn(3, 4)
```

String getColumn(int rowNumber, String columnName)

- *rowNumber* is the index of a row in the result set.
- *columnName* is the name of a column of the result set.

The method returns the text of the specified column.

For example, if *rsx* is a Java variable of type `ResultSetXml`:

```
select rsx>>getColumn(3, "name")
```

void setColumn (int rowNumber, int columnNumber, newValue)

rowNumber and *columnNumber* are as described for `getColumn()`.

The method sets the text of the specified column to *newValue*.

For example, if *rsx* is a Java variable of type ResultSetXml:

```
select rsx = rsx>>setColumn(3, 4, "new value")
```

void setColumn (int rowNumber, String columnName, newValue)

rowNumber and *columnName* are as described for getColumn().

The method sets the text of the specified column to *newValue*.

For example, if *rsx* is a Java variable of type ResultSetXml:

```
select rsx = rsx>>setColumn(3, "name", "new value")
```

Boolean allString (int ColumnNumber, String compOp, String comparand)

- *columnNumber* is the index of a column of the result set.
- *compOp* is a SQL comparison operator (<, >, =, !=, <=, >=).
- *comparand* is a comparison value.

The method returns a value indicating whether the specified comparison is true for all rows of the result set.

For example, if *rsx* is a Java variable of type ResultSetXml:

```
if rsx>>allString(3, "<", "compare value")...
```

This condition is true if in the result set represented by *rsx*, for all rows, the value of column 3 is less than "compare value." This is a String comparison. Similar methods can be used for other datatypes.

Boolean someString (int columnNumber, String compOp, String comparand)

- *columnNumber* is the index of a column of the result set.
- *compOp* is a SQL comparison operator (<, >, =, !=, <=, >=).
- *comparand* is a comparison value.

The method returns a value indicating whether the specified comparison is true for some row of the result set.

For example, if *rsx* is a Java variable of type `ResultSetXml`:

```
if rxs>>someString(3, "<", "compare value") ...
```

This condition is true if in the result set represented by *rsx*, for some row, the value of column 3 is less than “compare value.”

A customizable example for different result sets

This section demonstrates how you can store XML documents or the data that they contain in an Adaptive Server database using the `ResultSet` class and its methods for handling result sets. You can customize the `ResultSet` class for your database application.

Contrast the `ResultSet` document type and the `Order` document type:

- The `Order` document type is a simplified example designed for a specific purchase-order application, and its Java methods are designed for a specific set of SQL tables for purchase order data. See “The `OrderXml` class for order documents” on page 143.
- The `ResultSet` document type is designed to accommodate many kinds of SQL result sets, and the Java methods designed for it include parameters to accommodate different kinds of SQL queries.

For this example, you create and work with XML `ResultSet` documents that contain the same data as the XML `Order` documents.

First, create the `orders` table and its data:

```
create table orders
(customer_id varchar(5) not null,
order_date datetime not null,
item_id varchar(5) not null,
quantity int not null,
unit smallint default 1)
insert into orders values ("123", "1999/05/07", "987", 5, 1)
insert into orders values ("123", "1999/05/07", "654", 3, 12)
insert into orders values ("123", "1999/05/07", "579", 1, 1)
```

Also, create the following SQL table to store complete XML `ResultSet` documents:

```
create table resultset_docs
(id char(5),
rs_doc xml.resultsets.ResultSetXml)
```

The ResultSet document type

ResultSet documents consist of ResultSetMetaData followed by ResultSetData, as shown in the following general form:

```
<?xml version="1.0"?>
<!DOCTYPE ResultSet SYSTEM 'ResultSet.dtd'>
<ResultSet>
  <ResultSetMetaData>
  ...
</ResultSetMetaData>
  <ResultSetData>
  ...
</ResultSetData>
</ResultSet>
```

The ResultSetMetaData portion of an XML ResultSet consists of the SQL metadata returned by the methods of the JDBC ResultSet class. The ResultSetMetaData for the example result set is:

```
<ResultSetMetaData
  getColumnCount="7">
  <ColumnMetaData
    getColumnDisplaySize="25"
    getColumnLabel="Date"
    getColumnName="Date"
    getColumnType="93"
    getPrecision="0"
    getScale="0"
    isAutoIncrement="false"
    isCurrency="false"
    isDefinitelyWritable="false"
    isNullable="false"
    isSigned="false" />
  <ColumnMetaData
    getColumnDisplaySize="5"
    getColumnLabel="CustomerId"
    getColumnName="CustomerId"
    getColumnType="12"
    getPrecision="0"
    getScale="0"
```



```
isAutoIncrement="false"
isCurrency="false"
isDefinitelyWritable="false"
isNullable="false"
    isSigned="false" />
<ColumnMetaData
    getColumnDisplaySize="50"
    getColumnLabel="CustomerName"
    getColumnName="CustomerName"
    getColumnType="12"
    getPrecision="0"
    getScale="0"
    isAutoIncrement="false"
    isCurrency="false"
    isDefinitelyWritable="false"
    isNullable="false"
        isSigned="false" />
<ColumnMetaData
    getColumnDisplaySize="5"
    getColumnLabel="ItemId"
    getColumnName="ItemId"
    getColumnType="12"
    getPrecision="0"
    getScale="0"
    isAutoIncrement="false"
    isCurrency="false"
    isDefinitelyWritable="false"
    isNullable="false"
        isSigned="false" />
<ColumnMetaData
    getColumnDisplaySize="20"
    getColumnLabel="ItemName"
    getColumnName="ItemName"
    getColumnType="12"
    getPrecision="0"
    getScale="0"
    isAutoIncrement="false"
    isCurrency="false"
    isDefinitelyWritable="false"
    isNullable="false"
        isSigned="false" />
<ColumnMetaData
    getColumnDisplaySize="11"
    getColumnLabel="Quantity"
    getColumnName="Quantity"
    getColumnType="4"
```

```
        getPrecision="0"
        getScale="0"
        isAutoIncrement="false"
        isCurrency="false"
        isDefinitelyWritable="false"
        isNullable="false"
        isSigned="true" />
    <ColumnMetaData
        getColumnDisplaySize="6"
        getColumnLabel="unit"
        getColumnName="unit"
        getColumnType="5"
        getPrecision="0"
        getScale="0"
        isAutoIncrement="false"
        isCurrency="false"
        isDefinitelyWritable="false"
        isNullable="false"
        isSigned="true" />
</ResultSetMetaData>
```

The names of the attributes of `ColumnMetaData` are simply the names of the methods of the `JDBC ResultSetMetaData` class, and the values of those attributes are the values returned by those methods.

The `ResultSetData` portion of an XML `ResultSet` document is a list of *Row* elements, each having a list of *Column* elements. The text value of a *Column* element is the value returned by the `JDBC getString()` method for the column. The `ResultSetData` for the example is:

```
<ResultSetData>
  <Row>
    <Column name="Date">1999-07-04 00:00:00.0</Column>
    <Column name="CustomerId">123</Column>
    <Column name="CustomerName">Acme Alpha</Column>
    <Column name="ItemId">987</Column>
    <Column name="ItemName">Coupler</Column>
    <Column name="Quantity">5</Column>
    <Column name="unit">1</Column>
  </Row>
  <Row>
    <Column name="Date">1999-07-04 00:00:00.0</Column>
    <Column name="CustomerId">123</Column>
    <Column name="CustomerName">Acme Alpha</Column>
    <Column name="ItemId">654</Column>
    <Column name="ItemName">Connector</Column>
    <Column name="Quantity">3</Column>
```

```

    <Column name="unit">12</Column>
  </Row>
  <Row>
    <Column name="Date">1999-07-04 00:00:00.0</Column>
    <Column name="CustomerId">123</Column>
    <Column name="CustomerName">Acme Alpha</Column>
    <Column name="ItemId">579</Column>
    <Column name="ItemName">Clasp</Column>
    <Column name="Quantity">1</Column>
    <Column name="unit">1</Column>
  </Row>
</ResultSetData>
</ResultSet>

```

The XML DTD for the ResultSetXml document type

The DTD for the XML ResultSet document type is:

```

<!ELEMENT ResultSet (ResultSetMetaData ,
    ResultSetData)>
<!ELEMENT ResultSetMetaData (ColumnMetaData)+>
<!ATTLIST ResultSetMetaData getColumnCount CDATA
    #IMPLIED>
<!ELEMENT ColumnMetaData EMPTY>
<!ATTLIST ColumnMetaData
    getCatalogName CDATA #IMPLIED
    getColumnDisplaySize CDATA #IMPLIED
    getColumnLabel CDATA #IMPLIED
    getColumnName CDATA #IMPLIED
    getColumnType CDATA #REQUIRED
    getColumnName CDATA #IMPLIED
    getPrecision CDATA #IMPLIED
    getScale CDATA #IMPLIED
    getSchemaName CDATA #IMPLIED
    getTableName CDATA #IMPLIED
    isAutoIncrement (true|false) #IMPLIED
    isCaseSensitive (true|false) #IMPLIED
    isCurrency (true|false) #IMPLIED
    isDefinitelyWritable (true|false) #IMPLIED
    isNullable (true|false) #IMPLIED
    isReadOnly (true|false) #IMPLIED
    isSearchable (true|false) #IMPLIED
    isSigned (true|false) #IMPLIED
    isWritable (true|false) #IMPLIED
    >
<!ELEMENT ResultSetData (Row)*>
<!ELEMENT Row (Column)+>
<!ELEMENT Column (#PCDATA)>
<!ATTLIST Column

```

```
    null (true | false) "false"  
    name CDATA #IMPLIED
```

Using the element storage technique

This section uses the orders table to illustrate mapping between SQL data and XML ResultSet documents.

In “Composing a ResultSet XML document from the SQL data” on page 170, we generate an XML ResultSet document from the SQL data, assuming that we are the originator of the XML ResultSet document. We use the resulting XML ResultSet document to describe the ResultSet DTD.

Composing a ResultSet XML document from the SQL data

You can use Java methods to evaluate a given query and generate an XML result set with the query’s data. This example uses a constructor method of the ResultSetXml class. For example:

```
new xml.resultset.ResultSetXml  
    ("select 1 as 'a', 2 as 'b', 3 ", "none",  
     "yes", "antibes:4000?user=sa");
```

The method uses internal JDBC operations to execute the argument query, and then constructs the XML ResultSet for the query’s data.

We can invoke this constructor in a client or in the Adaptive Server:

- If you invoke the constructor in a client, specify a server parameter that identifies the Adaptive Server to be used when evaluating the query. The query is evaluated in the Adaptive Server, but the XML document is assembled in the client.
- If you invoke the constructor in the Adaptive Server, specify a null value or *jdbc:default:connection* for the server. The query is evaluated in the current server and the XML document is assembled there.

Generating a ResultSet in the client

The `main()` method of the `OrderResultSetClient` class is invoked in a client environment. `main()` invokes the constructor of the `ResultSetXml` class to generate an XML `ResultSet`. The constructor executes the query, retrieves its metadata and data using JDBC `ResultSet` methods, and assembles an XML `ResultSet` document with the data.

```
import java.io.*;
import util.*;
public class OrderResultSetClient {
    public static void main (String[] args) {
        try{
            String orderQuery = "select order_date as Date,"
                + "c.customer_id as CustomerId, "
                + "customer_name as CustomerName, "
                + "o.item_id as ItemId, i.item_name as ItemName, "
                + "quantity as Quantity, o.unit as unit "
                + "from customers c, orders o, items i "
                + "where c.customer_id=o.customer_id and "
                + "o.item_id=i.item_id " ;
            xml.resultset.ResultSetXml rsx
                = new xml.resultset.ResultSetXml (orderQuery,
                    "none", "yes", "external",
                    "antibes:4000?user=sa");
            FileUtil.string2File("OrderResultSet.xml",rsx.getXmlText());
        } catch (Exception e) {
            System.out.println("Exception:");
            e.printStackTrace();
        }
    }
}
```

Generating a result set in Adaptive Server

The following SQL script invokes the constructor of the `ResultSetXml` class in a server environment:

```
declare @rsx xml.resultset.ResultSetXml
select @rsx = new xml.resultset.ResultSetXml
    ("select 1 as 'a', 2 as 'b', 3 ", "none", "yes", "");
insert into resultset_docs values ("1", @rsx)
```

Translating the XML ResultSet document in the client

The `main()` method of `ResultSetXml` is executed in a client environment. It copies the file *OrderResultSet.xml*, constructs a *ResultSetXml* object containing the contents of that file, and invokes the `toSqlScript()` method of that object to generate a SQL script that re-creates the data of the result set. The method stores the SQL script in the file *order-resultset-copy.sql*.

```
import java.io.*;
import jcs.util.*;
public class ResultSet2Sql{
    public static void main (String[] args) {
        try{
            String xml = FileUtil.file2String("OrderResultSet.xml");
            xml.resultset.ResultSetXml rsx
                = new xml.resultset.ResultSetXml(xml);
            String sqlScript
                = rsx.toSqlScript("orderresultset_copy", "col_", "no");
            FileUtil.string2File("order-resultset-copy.sql", sqlScript);
            ExecSql.statement(sqlScript, "antibes:4000?user=sa");
        } catch (Exception e) {
            System.out.println("Exception:");
            e.printStackTrace();
        }
    }
}
```

This is the SQL script generated by `ResultSet2Sql`.

```
set quoted_identifier on
create table orderresultset_copy (
    Date datetime not null ,
    CustomerId varchar (5) not null ,
    CustomerName varchar (50) not null ,
    ItemId varchar (5) not null ,
    ItemName varchar (20) not null ,
    Quantity integer not null ,
    unit smallint not null
)
insert into orderresultset_copy values (
    '1999-07-04 00:00:00.0', '123',
    'Acme Alpha', '987', 'Widget', 5, 1 )
insert into orderresultset_copy values (
    '1999-07-04 00:00:00.0', '123',
    'Acme Alpha', '654',
    'Medium connector', 3, 12 )
insert into orderresultset_copy values (
```

```
'1999-07-04 00:00:00.0', '123',  
'Acme Alpha', '579', 'Type 3 clasp', 1, 1 )
```

The SQL script includes the `quoted_identifier` on command for those cases where the generated SQL uses quoted identifiers.

Translating the XML ResultSet Document in Adaptive Server

The following SQL script invokes the `toSqlScript()` method in Adaptive Server and then creates and populates a table with a copy of the result set data.

```
declare @rsx xml.resultset.ResultSetXml  
select @rsx = rs_doc from resultset_docs where id=1  
select @script = @rsx>>toSqlScript("resultset_copy",  
    "column_", "no")  
declare @I integer  
select @I = util.ExecSql.statement(@script, "")
```

Using the document storage technique

This section shows examples of storing XML ResultSet documents in single SQL columns and techniques for referencing and updating the column elements.

Storing an XML ResultSet document in a SQL column

The following SQL script generates an XML ResultSet document and stores it in a table:

```
declare @query java.lang.StringBuffer  
select @query = new java.lang.StringBuffer()  
-- The following "appends" build up a SQL select statement in  
    the @query variable  
-- We use a StringBuffer, and the append method, so that the  
    @query can be as long as needed.  
select @query>>append("select order_date as Date,  
    c.customer_id as CustomerId, ")
```

```
select @query>>append("customer_name as CustomerName, ")
select @query>>append("o.item_id as ItemId, i.item_name as ItemName, ")
select @query>>append("quantity as Quantity, o.unit as unit " )
select @query>>append("from customers c, orders o, items i ")
select @query>>append("where c.customer_id=o.customer_id and"
    + "o.item_id=i.item_id ")
declare @rsx xml.resultset.ResultSetXml
select  @rsx = new xml.resultset.ResultSetXml
        (@query>>toString(), 'none', 'yes', '')
insert into resultset_docs values("1",  @rsx)
```

Accessing the columns of stored ResultSet documents

In “Storing an XML ResultSet document in a SQL column” on page 173 , you inserted a complete XML ResultSet document into the `rs_doc` column of the `resultset_docs` table. This section shows examples of using methods of the `ResultSetXml` class to reference and update a stored ResultSet.

A client-side call

The `main()` method of the `ResultSetElements` class is executed in a client environment. It copies the file *OrderResultSet.xml*, constructs a `ResultSetXml` document from it, and then accesses and updates the columns of the ResultSet.

```
import java.io.*;
import util.*;
public class ResultSetElements {
    public static void main ( String[] args) {
        try{
            String xml =
                FileUtil.file2String("OrderResultSet.xml");
            xml.resultset.ResultSetXml rsx
                = new xml.resultset.ResultSetXml(xml);
            // Get the columns containing customer and date info
            String cname = rsx.getColumn(0, "CustomerName");
            String cid   = rsx.getColumn(0, "CustomerId");
            String date  = rsx.getColumn(0, "Date");
            // Get the elements for item 1 (numbering from 0)
            String iName1 = rsx.getColumn(1, "ItemName");
            String iId1  = rsx.getColumn(1, "ItemId");
            String iQ1   = rsx.getColumn(1, "Quantity");
            String iU    = rsx.getColumn(1, "unit");
            System.out.println("\nBEFORE UPDATE: ");
            System.out.println("\n    "+date+ "    "+ cname + "    " +
```



```

        cid);
System.out.println("\n    "+ iName1+" "+iId1+" "
    + iQ1 + " " + iU + "\n");
// Set the elements for item 1 (numbering from 0)
rsx.setColumn(1, "ItemName", "Flange");
rsx.setColumn(1, "ItemId", "777");
rsx.setColumn(1, "Quantity", "3");
rsx.setColumn(1, "unit", "13");
// Get the updated elements for item 1 (numbering
    from 0) iName1 = rsx.getColumn(1, "ItemName");
iId1 = rsx.getColumn(1, "ItemId");
iQ1 = rsx.getColumn(1, "Quantity");
iU = rsx.getColumn(1, "unit");
System.out.println("\nAFTER UPDATE: ");
System.out.println("\n    "+date+ "    "+ cname + " " +
    cid);
System.out.println("\n    "+ iName1+" "+iId1+" "
    + iQ1 + " " + iU + "\n");

// Copy the updated document to another file
FileUtil.string2File("Order-updated.xml",
    rsx.getXmlText());
} catch (Exception e) {
System.out.println("Exception:");
e.printStackTrace();
}
}
}

```

The `FileUtil.string2File()` method stores the updated `ResultSet` in the file *Order-updated.xml*. The `ResultSetMetaData` of the updated document is unchanged. The updated `ResultSetData` of the document is as follows with new values in the second item.

```

<ResultSetData>
  <Row>
    <Column name="Date">1999-07-04 00:00:00.0</Column>
    <Column name="CustomerId">123</Column>
    <Column name="CustomerName">Acme Alpha</Column>
    <Column name="ItemId">987</Column>
    <Column name="ItemName">Widget</Column>
    <Column name="Quantity">5</Column>
    <Column name="unit">1</Column>
  </Row>
  <Row>
    <Column name="Date">1999-07-04 00:00:00.0</Column>
    <Column name="CustomerId">123</Column>

```

```
<Column name="CustomerName">Acme Alpha</Column>
<Column name="ItemId">777</Column>
<Column name="ItemName">Flange</Column>
<Column name="Quantity">3</Column>
<Column name="unit">13</Column>
</Row>
<Row>
  <Column name="Date">1999-07-04 00:00:00.0</Column>
  <Column name="CustomerId">123</Column>
  <Column name="CustomerName">Acme Alpha</Column>
  <Column name="ItemId">579</Column>
  <Column name="ItemName">Type 3 clasp</Column>
  <Column name="Quantity">1</Column>
  <Column name="unit">1</Column>
</Row>
</ResultSetData>
</ResultSet>
```

A server-side script

Using the SQL script in “Storing an XML ResultSet document in a SQL column” on page 173, you stored complete XML ResultSet documents in the `rs_doc` column of the `resultset_docs` table. The following SQL commands, executed in a server environment, reference and update the columns contained in those documents.

You can select columns by name or by number:

- Select the columns of row 1, specifying columns by name:

```
select rs_doc>>getColumn(1, "Date"),
       rs_doc>>getColumn(1, "CustomerId"),
       rs_doc>>getColumn(1, "CustomerName"),
       rs_doc>>getColumn(1, "ItemId"),
       rs_doc>>getColumn(1, "ItemName"),
       rs_doc>>getColumn(1, "Quantity"),
       rs_doc>>getColumn(1, "unit")
  from resultset_docs
```

- Select the columns of row 1, specifying columns by number:

```
select rs_doc>>getColumn(1, 0),
       rs_doc>>getColumn(1, 1),
       rs_doc>>getColumn(1, 2),
       rs_doc>>getColumn(1, 3),
       rs_doc>>getColumn(1, 4),
       rs_doc>>getColumn(1, 5),
```

```
rs_doc>>getColumn(1, 6)
from resultset_docs
```

Specify some nonexistent columns and rows. Those references return null values.

```
Select rs_doc>>getColumn(1, "itemid"),
       rs_doc>>getColumn(1, "xxx"),
       rs_doc>>getColumn(1, "Quantity"),
       rs_doc>>getColumn(99, "unit"),
       rs_doc>>getColumn(1, 876)
from resultset_docs
```

Update columns in the stored ResultSet document:

```
update resultset_docs
set rs_doc = rs_doc>>setColumn(1, "ItemName",
"Wrench")
where id= "1"
update resultset_docs
set rs_doc = rs_doc>>setColumn(1, "ItemId", "967")
where id="1"
update resultset_docs
set rs_doc = rs_doc>>setColumn(1, "unit", "6")
where id="1"
select rs_doc>>getColumn(1, "ItemName"),
       rs_doc>>getColumn(1, "ItemId"),
       rs_doc>>getColumn(1, "unit")
from resultset_docs
where id="1"
```

Quantified comparisons in stored ResultSet documents

ResultSetXml contains two methods, `allString()` and `someString()`, for quantified searches on columns of a ResultSetXML document. To illustrate these two methods, first create some example rows in the `order_results` table.

The `order_results` table has been initialized with one row, whose `id` = "1" and whose `rs_doc` column contains the original Order used in all examples.

The following statements copy that row twice, assigning `id` values of "2" and "3" to the new rows. The `order_results` table now has three rows, with `id` column values of "1," "2," and "3" and the original Order.

```
insert into resultset_docs(id, rs_doc)
select "2", rs_doc
from resultset_docs where id="1"
```

```
insert into resultset_docs (id, rs_doc) select "3", rs_doc
from resultset_docs where id="1"
```

The following statements modify the row with an id column value of “1” so that all three items have ItemIds of “100”, “110”, and “120”:

```
update resultset_docs
  set rs_doc = rs_doc>>setColumn(0, "ItemId", "100")
  where id="1"
update resultset_docs
  set rs_doc = rs_doc>>setColumn(1, "ItemId", "110")
  where id="1"
update resultset_docs
  set rs_doc = rs_doc>>setColumn(2, "ItemId", "120")
  where id="1"
```

The following update statement modifies the row with id = “3” so that its second item (from 0) has an ItemId of “999”:

```
update resultset_docs
  set rs_doc = rs_doc>>setColumn(2, "ItemId", "999")
  where id="3"
```

The following select statement displays the id column and the three ItemId values for each row:

```
select id, rs_doc>>getColumn(0, "ItemId"),
       rs_doc>>getColumn(1, "ItemId"),
       rs_doc>>getColumn(2, "ItemId")
from resultset_docs
```

The results of the select are:

1	100	110	120
2	987	654	579
3	987	654	999

Note the following:

- The row with id of “2” is the original Order data.
- The row with id of “1” has been modified so that *all* ItemIds for that row are less than “200.”
- The row with id of “3” has been modified so that *some* ItemId for that row is greater than or equal to “9999”

The following expresses these quantifications with the allString() and someString() methods:

```
select id, rs_doc>>allString(3, "<", "200") as "ALL test"
```

```

from resultset_docs
select id, rs_doc>>someString(3, ">=", "999") as "SOME test"
from resultset_docs
select id as "id for ALL test" from resultset_docs
where rs_doc>>allString(3, "<", "200")>>booleanValue() = 1
select id as "id for SOME test" from resultset_docs
where rs_doc>>someString(3, ">=", "999")>>booleanValue() = 1

```

The first two statements show the quantifier in the select list and give these results:

ID	"all" test	"some" test
1	true	false
2	false	false
3	false	true

The last two statements show the quantifier in the where clause and give these results:

- ID for "all" test = "3"
- ID for "some" test = "1"

In the examples with the quantifier method in the where clause, note that:

- The where clause examples compare the method results with an integer value of 1. SQL does not support the Boolean datatypes as a function value, but instead treats Boolean as equivalent to integer values 1 and 0, for true and false.
- The where clause examples use the `booleanValue()` method. The `allString()` and `someString()` methods return type `java.lang.Boolean`, which is not compatible with SQL integer. The Java `booleanValue()` method returns the simple Boolean value from the Boolean object, which is compatible with SQL integer. This behavior is a result of merging the SQL and Java type systems.

The quantifier methods return `java.lang.Boolean` instead of simply Java boolean so that they can return null when the column is out of range, which is consistent with the SQL treatment of out-of-range conditions.

The following statements show quantifier references that specify column 33, which does not exist in the data:

```

select id, rs_doc>>allString(33, "<", "200") as "ALL test"
from resultset_docs
select id as "id for ALL test" from resultset_docs
where rs_doc>>allString(33, "<", "200")>>booleanValue() = 1

```

ID	“all” test
1	NULL
2	NULL
3	NULL

The ID for the “all” test = (empty).

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	181
Using the debugger	182
A debugging tutorial	189

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 189 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 *\$SYBASE/\$SYBASE_ASE/sample/JavaSql/manual-examples/JDBCExamples.java*.

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 16 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/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 SQL 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 `selector()` method from the Breakpoint window. This line is:

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

- 2 In Interactive SQL, 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 *sql* 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 SQL 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.

Network Access Using java.net

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	195
java.net classes	196
Setting up java.net	196
Example usage	197
User notes	202
Where to go for help	202

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 197.

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 1.1 shows the java.net classes Sybase supports.

Table 11-1: 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 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:

```
public static void writeOut(InputStream fin)
    throws Exception{
    Socket socket = new Socket("localhost", 1718);
    OutputStream fout = new
    BufferedOutputStream(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.

```
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) throws
    IOException{
    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.

```
create table t(c1 text)
go
insert values into t1 ("samplestring")
go
select TestStream2File.writeOut(c1) from t
```

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

Using the MailTo class URL to mail a document

Mailing a document is a good example of using the URL class. Before you start, you must have your client connected to a mailer, 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:

```
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 the MailTo class 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. Before you start, your client must connect 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.

To use the following code sample, you must:

- Read the entire document into Adaptive Server memory.
- Create a new InputStream on the document in Adaptive Server memory.

For example:

```
public static InputStream url_test()
    throws Exception
{
    URL u = new URL("http://www.xxxxxx.com/");
    Reader in = new InputStreamReader(u.openStream());
    int n=0, off;
    char c[]=new char[50000];
    for(off=0; (off<c.length-512)
    &&(n=in.read(c,off,512))!=-1; off+=n)
    for(off=0; off < c.length; off++) {
        b[off]=(byte)c[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 (c1 text)
insert into t values (

    URLprocess.readURL)
select datalength(c1) from
-----
34136

```

Accessing an external document with XQL

You can access an external document using the Adaptive Server XQL query function, which both parses and queries XML documents.

Pass the XML document to the XQL parser as an `InputStream`. You can use the class `URLProcess` to pass the XML document to either the XQL parse method or the XQL query method.

The class `URLProcess` is available on

```
select xml.Xql.query("//ItemID",
    URLProcess.readURL
    ("http://www.myserver.com/xmltest.xml"))
```

- `$SYBASE/ASE-12_5/sample/JavaSql` for UNIX environments
- `%SYBASEASE-12_5\sampleJavaSql` for NT environments

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.

Where to go for help

Reference documents:

- Java Examples in a Nutshell: A Desktop Quick Reference. David Flanagan, O'Reilly 1997

- Java Network Programming: Complete guide to networking, streams, and distributed computing. Hughes, Shoffner, Hamner, Bellur, Manning 1997

These documents are printed; you can find many more Java documents on the java.sun.com Web site.

Reference Topics

This chapter presents information on several reference topics.

Topic	Page
Assignments	205
Allowed conversions	207
Transferring Java-SQL objects to clients	207
Supported Java API packages, classes, and methods	208
Invoking SQL from Java	211
Transact-SQL commands from Java methods	212
Datatype mapping between Java and SQL	217
Java-SQL identifiers	219
Java-SQL class and package names	220
Java-SQL column declarations	221
Java-SQL variable declarations	221
Java-SQL column references	222
Java-SQL member references	223
Java-SQL method calls	224

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 222) 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 217. 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 217. 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 31 for more information.

- If the client is a Java client that uses `jConnect 4.0` or later, the server transmits the object serialization to the client. This serialization is seamlessly deserialized by `jConnect` to yield a copy of the object.
- If the client is a `bcp` client:

- 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.

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 209 for a list of the unsupported classes in `java.lang`.

- `java.math`
- `java.net` – see Chapter 11, “Network Access Using `java.net`”
- `java.sql` – see “Unsupported `java.sql` methods and interfaces” on page 209 for a list of the unsupported methods and interfaces in `java.sql`.
- `java.text`
- `java.util`
- `java.util.zip`

Unsupported Java packages and classes

- `java.applet`
- `java.awt`
- `java.awt.datatransfer`
- `java.awt.event`
- `java.awt.image`
- `java.awt.peer`
- `java.beans`
- `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 specifications. It is described at <http://www.javasoft.com>. `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 12-1 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 12-1: 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.
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.

Command	Status
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.
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.

Command	Status
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 12-2 lists set command options and whether or not you can use them in Java methods.

Table 12-2: 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
session_authorization	Not supported
showplan	Supported
sort_resources	Not supported

set command option	Status
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 12-3 shows SQL-Java datatype mapping.

Table 12-3: Mapping SQL datatypes to Java types

SQL type	Java type
char	String
varchar	String
nchar	String
nvarchar	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
real	float
float	double
double precision	double
binary	byte[]
varbinary	byte[]
image	byte[]
datetime	java.sql.Timestamp
smalldatetime	java.sql.Timestamp

Table 12-4 shows Java-SQL datatype mapping.

Table 12-4: Mapping Java scalar types to SQL datatypes

Java scalar type	SQL type
boolean	bit
byte	tinyint
short	smallint
int	integer
long	integer
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 (<code>_</code>) symbol [alphabetic character arabic numeral underscore(<code>_</code>) symbol dollar (<code>\$</code>) 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 (<code>_</code>) symbol. Subsequent characters can include alphabetic characters (uppercase or lowercase), numbers, the dollar (<code>\$</code>) symbol, or the underscore (<code>_</code>) 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 `set 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 <p>For Java-SQL package names:</p> <ul style="list-style-type: none"> • If you specify a Java-SQL subpackage name, you must reference the subpackage name with its package name: <pre>java_sql_package_name.java_sql_subpackage_name</pre>

- 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 <i>Reference Manual</i> .

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	<code>java_sql_variable ::= @variable_name java_sql_class_name</code>

	<i>java_sql_parameter ::= @parameter_name 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: <ul style="list-style-type: none"> • 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> <i>member_reference</i> ::= <i>class_member_reference</i> <i>instance_member_reference</i> <i>class_member_reference</i> ::= <i>java_sql_class_name</i>.<i>method_name</i> <i>instance_member_reference</i> ::= <i>instance_expression</i>>><i>member_name</i> <i>instance_expression</i> ::= <i>column_reference</i> <i>variable_name</i> <i>parameter_name</i> <i>method_call</i> <i>member_reference</i> <i>member_name</i> ::= <i>field_name</i> <i>method_name</i> </pre>
Parameters	<p><i>member_reference</i></p> <p>An expression that describes a field or method of a class or object.</p> <p><i>class_member_reference</i></p> <p>An expression that describes a static method of a Java-SQL class.</p> <p><i>instance_member_reference</i></p> <p>An expression that describes a static or dynamic method or field of a Java-SQL class instance.</p> <p><i>java_sql_class_name</i></p> <p>A fully qualified name of a Java-SQL class in the current database.</p> <p><i>instance_expression</i></p> <p>An expression whose datatype is a Java-SQL class.</p> <p><i>member_name</i></p> <p>The name of a field or method of the class or class instance.</p>
Usage	<ul style="list-style-type: none"> 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. <p>Otherwise, the Java-SQL member reference has the null value.</p> <ul style="list-style-type: none"> 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	<pre><i>method_call</i> ::= <i>member_reference</i> ([<i>parameters</i>]) new <i>java_sql_class_name</i> ([<i>parameters</i>]) <i>parameters</i> ::= <i>parameter</i> [(, <i>parameter</i>)...] <i>parameter</i> ::= <i>expression</i></pre>
Parameters	<p><i>method_call</i></p> <p>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.</p> <p><i>member_reference</i></p> <p>A member reference that denotes a method.</p> <p><i>parameters</i></p> <p>The list of parameters to be passed to the method. If there are no parameters, include empty parentheses.</p>
Usage	<p>Method overloading</p> <ul style="list-style-type: none">• When there are methods with the same name in the same class or instance, the issue is resolved according to Java method overloading rules. <p>Datatype of method calls</p> <ul style="list-style-type: none">• 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 select, fetch, insert, and update 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.
document type declaration (DTD)	In XML, every valid document has a DTD that describes the elements available in that document type. A DTD can be embedded in the XML document or referenced by it.

eXtensible Markup Language (XML)	A metalanguage designed for Web applications that lets you define your own markup tags and attributes for different kinds of documents. XML is a subset of SGML.
eXtensible Query Language (XQL)	A markup language for querying XML documents stored in a relational database. Adaptive Server provides an XQL query engine that can be installed in Adaptive Server or run as a standalone program
eXtensible Style Language (XSL)	A markup language designed to format XML documents into HTML or other XML documents with different attributes and tags.
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 .
Hypertext Markup Language (HTML)	A subset of SGML designed for the Web.
installed classes	Java classes and methods that have been placed in the Adaptive Server system by the installjava 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 217.
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 12-3 on page 218, 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 12-4 on page 219, 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	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.
narrowing conversion	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 .

package	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 import statements to specify other packages whose classes can then be referenced.
procedure	An SQL stored procedure, or a Java method with a <i>void</i> result type.
public	Public fields and methods, as defined in Java.
retained JAR	See associated JAR .
serialization	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 .
SQL function signature	The SQL datatype of each parameter of a SQLJ function.
SQL-Java datatype mapping	Conversions between Java and SQL datatypes. See “Datatype mapping between Java and SQL” on page 217.
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.
valid document	In XML, a valid document has a DTD and adheres to it. It is also a well-formed document .

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 convert function. See also narrowing conversion .

Index

Symbols

- , (comma)
 - in SQL statements xvii
- () (parentheses)
 - in SQL statements xvii
- [] (square brackets)
 - in SQL statements xvii
- >> (double angle)
 - to qualify Java fields and methods 223
- @ sign 82
- { } (curly braces) in SQL statements xvii

A

- access, server, to order elements 158
- Adaptive Server
 - installing XQL 122
 - plug-in 27, 80
- additional information
 - about Java 10
 - about XML 110
- ADT mappable datatypes 98
- allString, Java method 177
- alter table
 - command 27, 123
 - syntax 27
- ANSI standards 4
- appendItem, Java method 159
- aseutils methods, com.sybase.xml.xql.Xql
 - methods, specific to 139
- assignment properties
 - Java-SQL data items 32
- assignments 205
- attaching to a Java VM 183
- attributes, embedded in element tags 113

B

- bookstore.xml
 - authors example 131
 - DTD conforming 132
 - filename 131
 - validate command 132
 - Web page 133
 - XML example 131
- Boolean allString, Java method 164
- Boolean expressions, within filter operators 127
- Boolean someString, Java method 164
- breaking
 - on a class method 187
 - on a line number 186
 - using conditions 187
 - using counts 187
 - when execution is not interrupted 188
- breakpoints 186
- building parse tree with XML parser 119
- bypassing character-set conversions 118

C

- called on null input parameter 82
- case expressions 37, 86
- character encoding. See character sets
- character sets
 - Adaptive server plug-in 80
 - client server 118
 - conversions, bypassing 118
 - declared matching actual 113
 - declared, actual 118
 - default UTF8 113
 - specifying 112
 - specifying with SAX 119
 - translations, bypassed 113
 - unicode 27, 36, 80
 - XML 113, 118

- XML data 118
- character-set value, UTF8 default 118
- character-string operations in SQL 119
- charindex command 119
- child operator 125
- class names 220
- class subtypes 36–38
- classes. *See* Java classes
- CLASSPATH
 - environment variable 122
 - environment variables for UNIX and NT 122
 - standalone program 122
 - xerces.jar, xml.zip, runtime.zip 122
- clients
 - bcp 207
 - isql 207
- client-side JDBC 7
- code samples
 - HTML, Order example 114
 - XML, Info example 113
 - XML, Item example 114
- column
 - declarations 221
 - referencing 222
- column datatypes, requirements 25
- column declarations 221
- column references 222
- com 140
- com.sybase.xml.xql.store methods 140
- com.sybase.xml.xql.store.SybMemXmlStream, XQL interface 140
- com.sybase.xml.xql.Xql
 - methods, specific to 136, 137, 138, 140, 141
- com.sybase.xml.xql.XqlDriver
 - command 121
 - local files 130
 - querying XML documents 130
 - standalone program 130
 - syntax 130
 - using 130
- comma (,) in SQL statements xvii
- command main method 102
- commands
 - alter table 123
 - charindex 119
 - com.sybase.xml.xql.XqlDriver 121
 - create table 26, 27
 - debug 130
 - drop function 86
 - FileInputStream() 133
 - help 130
 - infile 130
 - insert 118
 - insert values clause of 118
 - new 122
 - outfile 130
 - parse 133
 - parse() 123
 - patindex 119
 - qstring 130
 - query 133
 - query() 128
 - select 128, 129
 - SQLJ create function 81
 - SQLJ create procedure 87
 - substring 119
 - update 122
 - URL 133
 - valid 132
 - validate 130, 132
 - where clause 128
 - writetext 118
- commands, create procedure SQLJ 90
- compile-time datatypes 38
- compiling Java code 16
- composing order documents 151
- configuration parameter, Number of Java Sockets 202
- configuring memory requirements 124
- constructor method 28
- constructors 28, 43
 - Order Xml 154
 - OrderXML 152
- conventions
 - Java-SQL syntax xv
 - Transact-SQL syntax xvi
- conversions 207
 - narrowing 37
 - widening 36
- convert function 36, 207
- create procedure (SQLJ) command 87, 90
- create table command, syntax 26, 27
- creating

- and populating SQL tables 149
- client applications 195
- network applications, java.net 195
- tables 26
- user-defined classes 16
- XSL style sheets 118
- curly braces ({}) in SQL statements xvii
- customizing elements 113

D

- data
 - selecting with XQL 121
 - translating from XML 153
- database objects
 - xmlcol 123, 125
 - xmlimage 125
 - XMLTEXT 123
 - XQL, general query language for XML 125
- DatagramPacket, Java class 196
- datatype conversions 207
- datatype mapping 35, 77, 98, 217–219
- datatypes
 - char 111
 - compile-time 38
 - conversions 207
 - image 111, 123, 124
 - Java classes 3
 - method calls 224
 - runtime 38
 - text 111, 123, 124
 - varchar 111
- debug command 130
- Debug.jar, Java file 182
- debugger
 - attaching to a Java VM 183
 - compiling classes for 183
 - disconnecting 188
 - how it works 181
 - location 182
 - options 185
 - requirements for using 181
 - starting 182
 - wait mode 183
- debugger capabilities
 - browse classes 182
 - inspect and break on expressions 182
 - inspect and set variables 182
 - set break conditions 182
 - set breakpoints 182
 - trace execution 182
- debugger windows
 - breakpoints 184
 - calls 184
 - classes 184
 - connection 184
 - exceptions 184
 - inspection 184
 - locals 185
 - source 184
- debugging
 - Java 181–193
- debugging tutorial 189–193
 - attaching to a Java VM 190
 - examples 191
 - inspecting local variables 192
 - inspecting static variables 193
 - inspecting variables 191
 - loading source code 190
 - modifying local variables 193
 - source code 189
 - starting the debugger 189
 - stepping through source code 191
- deleting 28, 97
 - Java objects 28
 - XML documents 124
- delimited identifiers 219
- descendant operator 125
- deterministic parameter 82, 88
- disabling Java 15
- distinct keyword 46
- Document Object Model (DOM) 119
- document storage 148, 149, 154, 159, 173–180
 - entire document 149
 - from client and server 154
 - tables 151
- Document Type Definition (DTD) 114
- Document Type Definition. See DTD
- document, ResultSet type 166
- document, validating 132
- DOM

- assembling parse tree 119
- building document parse tree 119
- generating parse tree 119
- modifying document parse tree 119
- object returned by SAX 119
- portable across XML parsers 119
- standard XML interface 119
- DOM, Document Object Model 119
- DOM. See Document Object Model
- double angle
 - qualifying Java fields and methods 223
 - to qualify Java fields and methods 29
- downloading
 - installed classes 21
 - installed JARs 21
- drop function command 86
- DTD 116
 - elements of 116
 - for valid XML document 117
 - internal 117
 - not required in all documents 117
- DTD, #IMPLIED 116
- DTD, #PCDATA DTD elements 116
- DTD, asterisk (*) 116
- DTD, ATTLIST 116
- DTD, document type definition 114
- DTD, ELEMENT 116
- DTD, plus sign (+) 116
- DTD, question mark (?) 116
- dynamic result sets parameter 88

E

- element storage 148, 149, 151, 154
 - 170–173
- extract elements and store 149
- element tags
 - customizing 113
 - embedded attributes 113
 - HTML, inconsistent 115
 - strict nesting 113
 - user-created 113
- elements
 - extracting 153
 - referencing and updating 155

- email
 - java.net 195
 - messages, sending 195
- embedding DTD in XML 117
- enabling Java 15
- enabling java.net, procedure 196
- equality operations 45
- equals, operator 126
- examples
 - for SQLJ routines 78
- exceptions 31
- explicit Java method signatures 99
- eXtensible Markup Language (XML). See Extensible Markup Language
- Extensible Markup Language. See XML
- Extensible Style Language. See XSL
- external name parameter 88
- external server, writing with java.net 197
- externalization 221
- extracting elements 153
- extractjava utility 21

F

- FileInputStream(), command 133
- filename, bookstore.xml 131
- filter operator 126
- filter operator, using Boolean expressions 127
- flushing data explicitly 202
- formatting
 - instructions, provided in XSL 111
 - XML information with XSL 118

G

- generating
 - document text from parse tree 119
 - Java parse tree 119
- generating order on client 152
- getString, JDBC method 168
- group by clause 46

H

- help command 130
- HoldString class 129
- HTML
 - display of Order data 114
 - DTD elements 116, 118
 - element bracketing inconsistent 115
 - javadoc-generated pages 110
 - Order code sample 114
 - subset of SGML. See Standardized General Markup Language
- HTML, Hypertext Markup Language 109
- URLConnection, Java class 196
- hybrid
 - storage tables 151
- hybrid storage 148, 149, 159–160
 - store document in XQL column 149
 - technique, using 159
- Hypertext Markup Language (HTML) 109

I

- identifiers 219
 - delimited 219
- image, datatype 123, 124
- implicit Java method signatures 99
- in parameter 90
- inconsistent element tags, HTML 115
- InetAddress, Java class 196
- infile, command 130
- info, XML code example 113
- inout parameter 90
- input sources, specifying with SAX 119
- InputStream class 200
- InputStream, Java class 201, 202
- insert command, values clause of 118
- inserting
 - data in a table 200
 - Java objects 28
 - XML document into database 118
- installing
 - compressed JARS 17
 - Java classes 17, 20
 - uncompressed JARS 17
 - XQL in Adaptive Server 122

- installjava utility 14, 17, 122
 - f option 18
 - j option 18
 - new option 19
 - syntax 18
 - update option 19
- instance methods 43
- inter-class arguments 51
- invoking
 - Java method, using SQLJ 79
 - Java methods 30, 78
 - Java methods, invoking directly 78
 - Java methods, using SQLJ 78
 - SQL from Java 211, 217

J

- JAR files
 - creating 17
 - installing 16
 - retaining 18
- JARs
 - compressed, installing 17
 - uncompressed, installing 17
- Java API 8
 - accessing from SQL 8
 - supported packages 208–211
 - Sybase support for 8
- Java arrays 91
- Java class datatypes 84
- Java classes 143
 - as datatypes 3, 25
 - creating 16
 - DatagramPacket 196
 - DatagramSocket 196
 - HoldString 129
 - URLConnection
 - InetAddress
 - InputStream 197, 200, 202
 - installing 17–20
 - JXm 161
 - JXml 144
 - MailTo 200
 - MulticastSocket 196
 - Order 159

- OrderXML 143
- OrderXml 143, 144, 161
- OutputStream 197, 200
- PrintWriter 202
- referencing other classes 20
- ResultSet 165
- ResultSetXml 161
- retained 21
- runtime 14
- saving in JAR 16
- ServerSocket 196, 198
- Socket
- SQLJ examples 78
- subtypes 36
- supported 8
- updating 19
- URL 200, 201
- URL class, using 198
- URLConnection 196
- URLDecoder 196
- URLEncoder 196
- user-defined 9, 14
- Java code
 - compiling 16
 - writing 16
- Java commands. See commands
- Java compiler 183
- Java constructor
 - OrderXml 144
- Java constructor ResultSetXm 162
- Java datatypes
 - ADT mappable 98
 - object mappable 98
 - output mappable 98
 - result-set mappable 99
 - simply mappable 98
- Java Development Kit 6
- Java in the database
 - advantages of 1
 - capabilities 2
 - key features 5
 - preparing for 13–22
 - questions and answers 5
- Java instances, representing 32
- Java method signature 83, 88
- Java methods
 - 177
 - allString 177
 - Boolean allString 164
 - Boolean someString 164
 - call by reference 31, 46
 - command main 102
 - constructor ResultSetXml 162
 - exceptions 31
 - getString, JDBC 168
 - instance 43
 - invoking 30, 78
 - OrderXml 144
 - ResultXml example 162
 - static 45
 - static void createOrderTable 145
 - String getColumn 163
 - String getItemElement 145
 - String toSqlScript 162
 - toSqlScript() 172
 - type 42, 43
 - void 43
 - void appendItem 147
 - void deleteItem(int itemNumber) 147
 - void order2Sql(String ordersTableName, String server) 144
 - void setColumn 163, 164
 - void setItemElement 146
 - void setOrderElement 145
 - XQL 136
- java methods
 - See also XQL methods
- Java methods, specific to com.sybase.xml.xql.Xql 136
- Java objects 28
- Java operations, invoked from SQL 7
- Java primitive datatypes 84
- Java runtime environment 13
- Java Services
 - increasing default memory parameters 121
 - memory requirements table 124
 - table, memory parameters 125
- Java VM 7, 13
- Java VM parameters
 - size of global fixed heap 197
 - size of process object heap 197
 - size of shared class heap 197
- Java, SQL, using together 7

- java.net 196, 197, 198, 202
 - accessing documents using XML, JDBC 195
 - accessing external documents 197
 - cautions 202
 - classes
 - client application, setting up 197
 - client process 198
 - client process procedure 198
 - connecting through JDBC with jconnect 195
 - creating networking applications 195
 - downloading documents 195
 - enabling 196
 - examples 197
 - help 202
 - mailing documents 197
 - objects not serializable 202
 - procedure for enabling 196
 - reference documents 202
 - references, online 202
 - references, written 202
 - saving documents 195
 - saving text from Adaptive Server 197
 - sending email messages 195
 - server process 198
 - server process procedure 198
 - writing external server 197
- java.net classes
 - HttpURLConnection 196
 - InetAddress 196
 - See Java classes
 - Socket 196
 - URL 196
 - URLConnection 196
 - URLDecoder 196
 - URLEncoder 196
- java.net, for network access 195
- java.sql 211
- java.sql methods, unsupported 209
- javadoc, generating HTML 110
- Java-SQL
 - class names 220
 - column declarations 221
 - column references 222
 - columns 33, 47
 - creating tables 26
 - function results 33
 - identifiers 219
 - member references 223
 - method calls 224
 - names 24
 - package names 220
 - parameters 33, 47
 - static variables 48
 - transferring objects 207
 - transferring objects to clients 207
 - unsupported methods 209
 - variable declarations 221
 - variables 33, 47
- Java-SQL classes
 - in multiple databases 48
 - installing 17–20
- Java-SQL columns
 - storage options 26
- jConnect
 - JDBC 7
 - OrderXml in client 152
- jconnect 195
- jConnect, used by OrderXml in client 152
- JDBC 57–74
 - accessing data 60
 - client-side 7, 58
 - concepts 58
 - connection defaults 59
 - connections 62
 - interface 9
 - JDBCExamples class 60
 - obtaining a connection 62
 - permissions 59
 - ResultSet class 165
 - ResultSetMetaData class 168
 - server-side 7, 58
 - terminology 58
 - version support 14
- JDBC drivers 14, 211
 - client-side 7, 58
 - jConnect 7
 - server-side 7, 58
 - used by OrderXml in server 152
- JDBC standard datatype mapping 98
- JDBCExamples class 69–74
 - methods 61–66
 - overview 60

L

language java parameter 88
local files, com.sybase.xml.xql.XqlDriver 130

M

mailing a document 197
MailTo, Java class 200
main() method, executed on client 156
mapping datatypes 217–219
mapping Java and SQL datatypes 98
mapping, illustrating with orders table 170
member references 223
memory parameters, Java Services, table 125
memory requirements
 configurng 124
 for query engine 124
 Java Services 124
memory requirements, Java Services parameters 121
method calls 224
 datatype of 224
method overloading 101, 224
methods
 appendItem 159
 exceptions 31
 main() executed on client 156
 order2Sql 153
 runtime results 225
 See also XQL methods
 SQLJExamples.bestTwoEmps() 78
 SQLJExamples.correctStates() 78, 89
 SQLJExamples.job() 78
 SQLJExamples.region() 78
 to reference and update elements 155
methods. See Java methods
modifies sql data parameter 82, 88
MulticastSocket, Java class 196
multiple databases 49

N

names in Java-SQL 24
 case 25
 length 24

narrowing conversions 37
navigating XQL 125
network access, java.net 195
new, command 122
null values
 case statements 86
 in SQLJ functions 84
nulls in Java-SQL 38–42
 arguments to methods 40
 using convert functions 41
Number of Java Sockets, configuration parameter 202

O

object mappable datatypes 98
obtaining connections 62
Open Client CT-Library 118
Open Client DB-Library 118
operators
 child 125
 descendant 125
 equal 126
 filter 126
 subscript 127
options
 external name 83
 language java 83
 parameter style java 83
 saxparser 131
options, storage, advantages and disadvantages 148
order
 generating on client 152
 generating on server 152
order by clauses 46
Order DTD, sample code 116
order elements, server access to 158
Order sample
 HTML 114
 XML code 111
Order, Java class 159
order2Sql, method 153
ordering operations 45
orders table 165
orders table, using 170
OrderXML

- Java class 143
- OrderXml 143
 - class 143
 - constructor invoked from server 154
 - Java class 143, 161
 - sample application 143
 - sample class 143
 - source code 143
 - subclass of JXml class 144
- OrderXml, Java method 144
- out parameter 90
- outfile, command 130
- output HTML, from XSL 118
- output mappable datatypes 98

P

- package names 220
- parameter style java parameter 88
- parameters
 - 130
 - (Java VM) size of global fixed heap 197
 - (Java VM) size of process object heap 197
 - (Java VM) size of shared class heap 197
 - deterministic 88
 - external name 88
 - help 130
 - inout 90
 - input 90
 - language java 88
 - modifies sql data 88
 - not deterministic 88
 - output 90
 - parameter style java 88
 - qstring 130
 - validate 130
- parentheses ()
 - in SQL statements xvii
- parse command 133
- parse methods 136
- parse tree
 - assembling with DOM 119
 - build or modify 119
 - generating document's text from 119
 - generating Java representation 119

- parse tree, building and modifying with DOM 119
- parse()
 - command 123
 - returns sybase.aseutils.SybXmlStream 123
- parse(), Java method, command 123
- parse(InputStream xmlIStream), XML method 137
- parse(String xmlDoc), XQL method 136
- parser, XML 121
- parsers for XML 119
- parsers, XML 119
- patindex command 119
- permissions
 - Java 7, 24
 - JDBC 59
 - SQLJ routines 77
- persistent data items 33
- plus sign (+) in XML document type definition 116
- presentation applications, use of XSL for 118
- PrintWriter, Java class 202
- procedure
 - creating SQLJ routine 76
 - enabling java.net 196
- procedures
 - client process, java.net 198
 - server process, java.net 198
- processing
 - effect of where clause 128
- processing with SAX, incremental 119
- processing XML, specialized 143

Q

- qstring command 130
- query
 - command 133
- query engine
 - as standalone program 121
 - inside or outside server 121
 - memory requirements 124
- query methods, com.sybase.xml.xql.Xql 138
 - methods, specific to 138, 139
- query structures 128
- query() command 128

query(String query, InputStream xmlDoc), XQL method 138
query(String query, String xmlDoc), XML method 138
query(String query, JXml jxml), XQL method 139
query(String query, SybXmlStream xmlDoc), XQL method 139
querying XML with com.sybase.xml.xql.XqlDriver 130
questions and answers 5

R

reading XML from database 118
rearranging installed classes 21
referencing
 fields 29
 XML DTD externally 117
remove java command 21, 221
removing classes 21
removing JARs 21
restrictions on Java in the database 10
result sets 101
 not stored in where clause 129
 unexpected result 129

ResultSet
 accessing columns of stored documents 174
 class 161
 composing document from SQL 170
 document type 166
 DTD 169
 generating in Adaptive Server 171
 generating in client 171
 Java class 165
 JDBC class 166
 mappable datatypes 98
 quantified comparisons in stored documents 177
 quantifier in select list 179
 quantifier in where clause 179
 search methods 177
 selecting and updating columns 176–177
 server-side script 176
 storing document in SQL column 173
 translating in client 172
ResultSetData 165
ResultSetMetaData 165
ResultSetMetaData for example result 166

ResultSetMetaData, JDBC class 168
ResultSetXml 161
 accessing XML 161
 directories, XML 161
 processing SQL result sets 161
 similar to OrderXml class 161
 source code 161
 source code, XML 161
 subclass of JXml class 161
 subset of JXml class 161
 writing Java code to access XML 161
ResultSetXml(String), Java method 162
ResultXml example, Java method 162
returns null on null input parameter, Java clause 82
runtime
 datatypes 38
Runtime environment 13
Runtime Java classes
 location of 14
runtime Java classes 14

S

sample classes 52–55
 address 52
 address2Line 53
 JDBCExamples 60–74
 location of 11
 misc 55
 OrderXml 143
 ResultSet 165
sample code
 DTD, Order example 116
 HTML, Order sample 114
 XML, Info example 113
 XML, Order example DTD 116
saving text out of Adaptive server 197
SAX
 generating events 119
 portable across XML parsers 119
 returning DOM object 119
 standard XML interface 119
SAX (Simple API for XML) 119
saxparser, com.sybase.xml.xql.XqlDriver option 131
search order

- function types 84
- searching, XML documents stored on Web 110
- security
 - SQLJ routines 77
- select command 128, 129
- selecting data with XQL 121
- selecting Java objects 28
- serialization 221, 222
- server access to order elements 158
- server process 198
- server-side JDBC 7
- ServerSocket, Java class 196, 198
- set commands
 - allowed in Java methods 216
 - updating 44
- setParser, XQL method 140
- setting up 196
- SGML, Standardized General Markup Language 109
- shared class heap 196
- simple API for XML (SAX) 119
- simply mappable datatypes 98
- Socket classes, using 197
- Socket, Java class 196
- someString, Java method 177
- sp_configure system procedure 15
- sp_depends system procedure 97
- sp_help system procedure 97
- sp_helpjava
 - syntax 20
 - utilitysp_helpjava 20
- sp_helpjava system procedure 98
- sp_helpprotect system procedure 98
- specialized XML processing 143
- specifying character set 112
- SQL
 - character-string operations 119
 - expressions, include Java objects 7
 - function signature 82
 - procedure signature 88
 - tables, creating and populating 149
 - 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 104
- SQLJ standards 76
- SQLJ stored procedures 87–89, 97
 - capabilities of 87
 - deleting 97
 - modifying SQL data 89
 - 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
- standalone applications
 - example 133
 - using XQL 133
- standalone program, com.sybase.xml.xml.XqlDriver 130
- Standardized General Markup Language (SGML) 109
- standards specifications 4
- static methods 45, 77, 79, 87
- static variables 48
- static void createOrderTable, Java method 145
- storage
 - document 149
 - element 149
 - hybrid 149
- storage options
 - advantages and disadvantages 148
 - in row 26
- storage tables, document and hybrid 151
- storage technique, hybrid, using 159
- storage, document 154
- storage, element 149
- storage, hybrid 149
- storing XML documents 147
- parameters
 - infile 130
 - outfile 130
- String data
 - zero length 42

Index

- string data 42
- String getColumn, Java method 163
- String getItemElement, Java method 145
- String toSqlScript, Java method 162
- structures, query 128
- style java keyword 88
- style sheets, creating 118
- style sheets, XSL 118
- subscript operator 127
- substring command 119
- subtypes 36
- supertypes 36
- Sybase Central
 - creating a SQLJ function or procedure from 80
 - managing SQLJ procedures and functions from 80
 - viewing SQLJ routine properties from 81
- sybase.asciutils 139
- sybase.aseutils.SybXmlStream, returned by parse()
 - command 123
- SybFileXmlStream, XQL method 140
- SybXmlStream 136
 - variable 136
- syntax conventions
 - Java-SQL xv
 - Transact-SQL xvi
- system procedures
 - helpjava 20
 - sp_depends 97
 - sp_help 97
 - sp_helpjava 98
 - sp_helpprotect 98

T

- table definition 78
- table, orders 165
- tags
 - customizing in XML 111
 - HTML, inconsistent bracketing 115
 - HTML, paragraph 115
 - user-created 113
 - XML strictly nested 112
- temporary databases 51
- text data, XML 113
- text, datatype 123, 124

- toSqlScript(), Java method 172
- transact-SQL
 - commands, in Java methods 212
- transient data items 33
- translating data 151, 153

U

- unicode 42
- union operator 46
- update, command 122
- updating Java objects 28
- updating XML documents 124
- URL
 - Java class 198
- URL class
 - accessing external server with XQL 200
 - downloading HTTP document 200
 - inserting data in a table 200
 - Java class 196, 200, 201
 - obtaining an HTTP document 200
 - sending email 200
 - using 200
- URL, command 133
- URLConnection, Java class 196
- URLDecoder, Java class 196
- URLEncoder, Java class 196
- user-created element tags 113
- user-defined classes, creating 16
- using
 - com.sybase.xml.xql.XqlDriver 130
 - hybrid storage technique 159
 - Java and SQL together 7
 - Java classes 23, 52
 - orders table 170
 - Socket classes 197
 - URL class 198
- UTF8, default character set 113, 118

V

- valid XML document 117
- valid, command 132
- validate, command 130, 132

- validating document 132
- variable declarations 221
- variables 221
 - datatypes of 25
 - static 48
 - SybXmlStream 136
 - values assigned to 28
- variables, HTML
 - ... 115
 - <table>...</table>, layout 115
 - bcolor, color 115
 - CustomerID 114, 115
 - CustomerName 115
 - Data 115
 - ItemID 114
 - ItemName 114
 - order 114
 - Quantities 115
 - Quantity 114
 - units 114
- variables, XML
 - OrderXML 110
 - tag 112
 - XMLResultSet 110
- viewing information
 - about installed classes 20
 - about installed JARs 20
- void appendItem, Java method 147
- void deleteItem(int itemNumber), Java method 147
- void methods 89
- void order2Sql(String ordersTableName, String server), Java method 144
- void setColumn, Java method 163, 164
- void setItemElement, java method 146
- void setOrderElement, Java method 145

W

- Web addresses
 - W3C, Document Object Model (DOM) 110
 - W3C, Extensible Markup Language (XML) 110
 - W3C, Extensible Stylesheet Language (XSL) 110
 - World Wide Web Consortium (W3C) 110
- Web information on XML 110

- Web, storage for XML documents 133
- well-formed XML document 113
- where clause 36, 44, 47
 - affects processing 128
 - command 128
 - not for result set storage 129
- work databases 51
- writetext, command 118

X

- xerces.jar
 - directory 123
- XML 109
 - accessing documents with java.net 195
 - additional information 110
 - application-specific document types 111
 - character sets, client and server 118
 - compared to SGML and HTML 111
 - comparison with HTML 109
 - custom tags 111
 - customizable example 165
 - declaring character sets 118
 - DTD (document type definition) 114
 - DTD elements, restrictions 116, 118
 - DTD not required in all documents 117
 - DTD sample code, embedding 117
 - DTD sample code, referencing externally 117
 - DTD, instruction 116, 118
 - Extensible Markup Language
 - input from XSL 118
 - overview 110, 111
 - parser, outside or inside server 121
 - parsers 119
 - read by HTML browsers and processors 111
 - sample document 111
 - source code for sample classes 110
 - specialized processing 143
 - storage options, pros and cons 148
 - strict phrase structure 111
 - subset of SGML 109
 - suitable for data interchange 109
 - tools written in Java 110
 - using XSL to transform 118
 - Web documents for detailed information 110

- XML data
 - document storage document storage, XML data 148
 - element storage 148
 - inserting from client file 155
- XML data operations
 - server-side 149
- XML declaration, to specify character set 112
- XML document
 - accessing elements 155
 - accessing in XQL 147
 - as character data 112
 - character-set value 118
 - client or server 147
 - deleting 124
 - DTD sample code 116
 - example, info 113
 - executing on server 155
 - formatting for 118
 - generating from Adaptive Server 110
 - generating from SQL data 109
 - inserting into database 118
 - mapping and storage 147
 - nested markup tags 112
 - no formatting instructions 113
 - parts of 112
 - querying 130
 - reading from database 118
 - sample code, Order 111
 - search with XQL 110
 - searching on Web 110
 - stored as OS file 133
 - stored on Web 133
 - storing 147
 - storing in Adaptive Server 110
 - updating 124
 - valid with DTD 117
 - well-formed 113
- XML document types
 - Order 165
 - ResultSet 165
- XML example, bookstore.xml 131
- XML methods
 - parse(InputStream xml_document) 137
 - query(String query, String xmlDoc) 138
- XML operations
 - client-side 149
- XML parser 119
 - applications using SAX and DOM portable 119
 - free license 119
 - public domain 119
 - standard interface 119
 - standard interfaces 119
- XML parsing with SAX 119
- XML Query Language (XQL) 110
- xml.zip, directory 123
- xmlcol, database object 123, 125
- xmllimage, database object 125
- XLTEXT, database object 123
- XQL
 - developing standalone applications 133
 - displaying as XML 121
 - EJBs 133
 - installing in Adaptive Server 122
 - interface,
 - com.sybase.xml.xql.store.SybMemXmlStream 140
 - JavaBeans 133
 - JDBC clients 133
 - navigation 125
 - numbering system 131
 - operators 125
 - parse method 202
 - parses and queries 202
 - path-based query language 110, 125
 - query method 202
 - to search XML documents 110
 - zero based 131
- XQL methods 136
 - parse(String xmlDoc) 136
 - query(String query, InputStream xmlDoc) 138
 - query(String query, JXml jxml) 139
 - query(String query, SybXmlStream xmlDoc) 139
 - setParser 140
 - SybFileXmlStream 140
 - SybXmlStream 139
- XQL, XML Query Language 110
- XSL 118
 - Extensible Style Language 111
 - formatting XML information 118
 - specifications 118
 - transforms XML 118
 - using with presentation applications 118

Z

zero-length strings 42

