



Programmer's Reference

jConnect™ for JDBC™

5.5

DOCUMENT ID: DC39001-01-0550-04

LAST REVISED: September 2003

Copyright © 1989-2003 by Sybase, Inc. All rights reserved.

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

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

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

Sybase, the Sybase logo, AccelaTrade, 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, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, AvantGo, AvantGo Application Alerts, AvantGo Mobile Delivery, AvantGo Mobile Document Viewer, AvantGo Mobile Inspection, AvantGo Mobile Marketing Channel, AvantGo Mobile Pharma, AvantGo Mobile Sales, AvantGo Pylon, AvantGo Pylon Application Server, AvantGo Pylon Conduit, AvantGo Pylon PIM Server, AvantGo Pylon Pro, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convoy/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTF, eFulfillment Accelerator, 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, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, Mail Anywhere Studio, MainframeConnect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, My AvantGo, My AvantGo Media Channel, My AvantGo Mobile Marketing, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, 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 Net Library, PocketBuilder, Pocket PowerBuilder, 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, Rapport, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, 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, S.W.I.F.T. Message Format Libraries, 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, TradeForce, Transact-SQL, Translation Toolkit, UltraLite.NET, 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. 03/03

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., One Sybase Drive, Dublin, CA 94568.

Contents

About This Book	vii
 CHAPTER 1	
Introduction	1
What is JDBC?	1
What is jConnect?	2
 CHAPTER 2	
Programming Information	5
Setting up jConnect	5
Setting the jConnect version	6
Invoking the jConnect driver	11
Establishing a connection	11
Setting connection properties	11
Connecting to Adaptive Server	20
Connecting to Adaptive Server Anywhere	21
Connecting to a server using JNDI	22
Implementing custom socket plug-ins	28
SYB SOCKET_FACTORY connection property	29
Creating and configuring a custom socket	29
Handling internationalization and localization	32
Using jConnect to pass Unicode data	33
jConnect character-set converters	34
Working with databases	39
Implementing high availability failover support	40
Performing server-to-server remote procedure calls	45
Wide table support for Adaptive Server version 12.5 and later	46
Accessing database metadata	47
Using cursors with result sets	48
Support for batch updates	59
Updating a database from a result set of a stored procedure	61
Working with datatypes	62
Implementing advanced features	68
Using event notification	69
Handling error messages	71

Storing Java objects as column data in a table	76
Dynamic class loading	80
JDBC 2.0 optional package extensions support	84
Restrictions on and interpretations of JDBC standards	95
Connection.isClosed() and the IS_CLOSED_TEST connection property	95
Statement.close() with unprocessed results	96
Making adjustments for multithreading	97
Using ResultSet.getCursorName()	98
Using setLong() with large parameter values	98
Using COMPUTE statements	98
Executing stored procedures	99

CHAPTER 3	Troubleshooting	101
	Debugging with jConnect	101
	Obtaining an instance of the Debug class	101
	Turning on debugging in your application	102
	Turning off debugging in your application	102
	Setting the CLASSPATH for debugging	103
	Using the Debug methods	103
	Capturing TDS communication	104
	PROTOCOL_CAPTURE connection property	105
	pause() and resume() methods in the Capture class	105
	Unsuccessful connection errors	106
	Gateway connection refused	106
	Unable to connect to a 4.9.2 SQL Server	107
	Memory usage in jConnect applications	107
	Stored procedure errors	108
	RPC returns fewer output parameters than registered	108
	Fetch/state errors when output params returned	108
	Stored procedure executed in unchained transaction mode ..	109
	Custom socket implementation error	109

CHAPTER 4	Performance and Tuning	111
	Improving jConnect performance	111
	BigDecimal rescaling	112
	REPEAT_READ connection property	112
	Character-set conversion	113
	Performance tuning for prepared statements in dynamic SQL	114
	Choosing prepared statements and stored procedures	115
	Prepared statements in portable applications	115
	Prepared statements with jConnect extensions	116
	Connection.prepareStatement()	117

	DYNAMIC_PREPARE connection property	118
	SybConnection.prepareStatement()	119
	ESCAPE_PROCESSING_DEFAULT connection property ...	120
	Cursor performance	120
	LANGUAGE_CURSOR connection property	121
CHAPTER 5	Migrating jConnect Applications	123
	Migrating jConnect applications	123
	Migrating applications to jConnect 4.5 and 5.5	123
	Sybase extensions	125
	Change example	125
	Method names.....	126
	Debug class.....	126
CHAPTER 6	Web Server Gateways	127
	About Web server gateways	127
	TDS tunnelling	127
	jConnect and gateway configuration	128
	Usage requirements	132
	Reading the index.html file	132
	Running the sample Isql applet	133
	Using the TDS-tunnelling servlet.....	134
	TDS-tunnelling servlet system requirements.....	135
	Installing the servlet.....	136
	Invoking the servlet	137
	Tracking active TDS sessions	137
	Resuming a TDS session.....	138
	TDS and Netscape Enterprise Server 3.5.1 on Solaris	138
APPENDIX A	SQL Exception and Warning Messages	141
APPENDIX B	jConnect Sample Programs	163
	Running IsqlApp	163
	Running jConnect sample programs and code	165
	Sample applications	165
	Sample code	166

About This Book

The *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.

Audience

This manual is for database-application programmers who are familiar with the Java programming language, JDBC, and Transact-SQL®, the Sybase version of Structured Query Language.

Related Documents

You may find the following documents helpful:

- The *Sybase jConnect for JDBC Installation Guide*
- The *Sybase jConnect for JDBC Release Bulletin*
- The javadoc documentation of jConnect extensions to JDBC. The Java Development Kit (JDK) from Java Software contains a javadoc script for extracting comments from source-code files. This script has been used to extract documentation of jConnect packages, classes, and methods from jConnect source files. When you install jConnect using the full installation or javadocs option, the javadoc information is placed in the *javadocs* directory:

Installation_directory/docs/en/javadocs

Other sources of information

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

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the Technical Library CD. It is included with your software. To read or print documents on the Getting Started CD you need Adobe Acrobat Reader (downloadable at no charge from the Adobe Web site, using a link provided on the CD).
- The Technical Library CD contains product manuals and is included with your software. The DynaText reader (included on the Technical Library CD) 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.

- The 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 EBFs/Updates, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

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.

❖ Finding 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 and click Go.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

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

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

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

Sybase EBFs and software updates

❖ Finding the latest information on EBFs and software updates

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 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 Select a product.

- 4 Specify a time frame and click Go.
- 5 Click the Info icon to display the EBF/Update report, or click the product description to download the software.

Conventions

This manual uses the following font and syntax conventions:

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

SybConnection class

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.”

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

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

- Code fragments are shown in a monospaced 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);
```

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.

Introduction

This chapter introduces you to jConnect for JDBC and describes its concepts and components.

This chapter contains:

Topics

What is JDBC?

What is jConnect?

What is JDBC?

JDBC (Java Database Connectivity) from the Java Software Division of Sun Microsystems, Inc. is a specification for an application program interface (API) that allows Java applications to access multiple database management systems using Structured Query Language (SQL). The JDBC driver manager handles multiple drivers that connect to different databases.

A set of interfaces is included in the standard JDBC API so you can open connections to databases, execute SQL commands, and process results. The interfaces are described in Table 1-1.

Table 1-1: JDBC interfaces

Interface	Description
java.sql.Driver	Locates the driver for a database URL
java.sql.Connection	Connection to a specific database
java.sql.Statement	Executes SQL statements
java.sql.PreparedStatement	Handles parameterized SQL statements
java.sql.CallableStatement	Handles database stored procedure calls
java.sql.ResultSet	Gets the results of SQL statements
java.sql.DatabaseMetaData	Used to access a variety of information about a connection's DBMS and database
java.sql.ResultSetMetaData	Used to access a variety of information describing a ResultSet's attributes

Each relational database management system requires a driver to implement these interfaces. All JDBC calls are sent to the JDBC driver manager, which passes the call to the specified driver.

There are four types of JDBC drivers:

- *Type 1 JDBC-ODBC bridge* – translates JDBC calls into ODBC calls and passes them to an ODBC driver. Some ODBC software must be resident on the client machine. Some client database code may also reside on the client machine.
- *Type 2 native-API partly-Java driver* – converts JDBC calls into database-specific calls. This driver, which communicates directly with the database server, also requires some binary code on the client machine.
- *Type 3 net-protocol all-Java driver* – communicates to a middle-tier server using a DBMS-independent net protocol. A middle-tier gateway then converts the request to a vendor-specific protocol.
- *Type 4 native-protocol all-Java driver* – converts JDBC calls to the vendor-specific DBMS protocol, allowing client applications direct communication with the database server.

What is jConnect?

jConnect is Sybase's high-performance JDBC driver. jConnect is both:

- A net-protocol/all-Java driver within a three-tier environment, and
- A native-protocol/all-Java driver within a two-tier environment.

The protocol used by jConnect is TDS 5.0 (Tabular Data Stream™, version 5), the native protocol for Adaptive Server® and Open Server™ applications. jConnect implements the JDBC standard to provide optimal connectivity to the complete family of Sybase products, allowing access to over 25 enterprise and legacy systems, including:

- Adaptive Server Enterprise
- Adaptive Server Anywhere
- Adaptive Server IQ (formerly Sybase IQ™)
- Replication Server®

- OmniConnect™

Note Since changing the name of Sybase SQL Server™ to Adaptive Server Enterprise, Sybase may use the names Adaptive Server and Adaptive Server Enterprise to refer collectively to all supported versions of Sybase SQL Server and Adaptive Server Enterprise. From this point forward, in this document, Adaptive Server Enterprise will be referred to as Adaptive Server.

In addition, jConnect for JDBC can access Oracle, AS/400, and other data sources using Sybase DirectConnect™.

In some instances, jConnect's implementation of JDBC deviates from the JDBC 1.x or 2.x specifications. For more information, see "Restrictions on and interpretations of JDBC standards" on page 95.

Programming Information

This chapter describes the basic components and programming requirements that comprise jConnect for JDBC. It explains how to invoke the jConnect driver, set connection properties, and connect to a database server. It also contains information about using jConnect features.

Note For information about JDBC programming, go to <http://java.sun.com/jdbc>.

To access the *JDBC Guide: Getting Started* manual for JDBC 1.0, go to <http://java.sun.com/products/jdk/1.1/docs/guide/jdbc>.

To access the *JDBC Guide: Getting Started* manual for JDBC 2.1, go to <http://java.sun.com/products/jdk/1.2/docs/guide/jdbc/>.

The following topics are included in this chapter:

Topics

Setting up jConnect

Establishing a connection

Implementing custom socket plug-ins

Handling internationalization and localization

Working with databases

Implementing advanced features

Restrictions on and interpretations of JDBC standards

Setting up jConnect

This section describes the tasks you need to perform before you use jConnect.

Setting the jConnect version

There are several versions of jConnect; use a version setting to determine:

- The default value of the LANGUAGE connection property.
- The version-specific features that are available.
- The default character set, if no character set is specified through the CHARSET connection property.
- The default value of the CHARSET_CONVERTER connection property.
- The default value of the CANCEL_ALL connection property, which is used to set the behavior of `Statement.cancel()`, which by default cancels the object on which it is invoked and any other `Statement` objects that have begun to execute and are waiting for results.
- If you are requesting support for wide tables from the server.
- If you would like to request server support for storing character data in `unichar` (Unicode) columns.

Note Only Adaptive Server version 12.5 and later support wide tables and `unichar` character data.

- If you would like to request support from the server for the date and time SQL datatypes.

Note Only Adaptive Server version 12.5.1 and later support the date and time SQL datatypes.

Table 2-1 lists the version settings available and their features.

Table 2-1: jConnect version settings and their features

Version constant	Features	Comments
VERSION_6	<ul style="list-style-type: none"> • jConnect requests support for the date and time SQL datatypes from the server. This request is ignored by servers other than Adaptive Server version 12.5.1 and later. • jConnect requests support for the unichar and univarchar datatypes from the server. This request is ignored by servers other than Adaptive Server 12.5 and later. • jConnect requests support for wide tables from the server. This request is ignored by servers other than Adaptive Server 12.5 and later. • The default value of the LANGUAGE connection property is null. • If the CHARSET connection property does not specify a character set, jConnect uses the database's default character set. The default value for CHARSET_CONVERTER is the PureConverter class. • By default, Statement.cancel() cancels only the Statement object it is invoked on. • JDBC 2.0 methods can be used to store and retrieve Java objects as column data. 	<p>For jConnect version 5.x, the default is VERSION_5.</p> <p>For additional information, see the comments for VERSION_4.</p> <p>For additional information for date and time datatypes, see "Using Date and Time datatypes" on page 67.</p> <p>For more information on wide tables, see "Wide table support for Adaptive Server version 12.5 and later" on page 46.</p> <p>For more information on unichar and univarchar datatypes and Unicode, see "Using jConnect to pass Unicode data" on page 33.</p>
VERSION_5	<ul style="list-style-type: none"> • The default value of the LANGUAGE connection property is null. • If the CHARSET connection property does not specify a character set, jConnect uses the database's default character set. The default value for CHARSET_CONVERTER is the PureConverter class. • By default, Statement.cancel() cancels only the Statement object it is invoked on. • JDBC 2.0 methods can be used to store and retrieve Java objects as column data. 	<p>For jConnect version 5.x, the default is VERSION_5.</p> <p>For additional information, see the comments for VERSION_4.</p>

Version constant	Features	Comments
VERSION_4	<ul style="list-style-type: none"> The default value of the LANGUAGE connection property is null. If the CHARSET connection property does not specify a character set, jConnect uses the database's default character set. The default value for CHARSET_CONVERTER is the PureConverter class. By default, Statement.cancel() cancels only the Statement object it is invoked on. JDBC 2.0 methods can be used to store and retrieve Java objects as column data. 	<p>Server messages are localized according to the language setting in your local environment. The languages supported are: Chinese, U.S. English, French, German, Japanese, Korean, Polish, Portuguese, and Spanish.</p> <p>The default behavior of Statement.cancel() is JDBC-compliant.</p> <p>Use CANCEL_ALL to set the behavior of Statement.cancel(). See "CANCEL_ALL connection property" on page 10.</p> <p>For information on Java objects as column data, see "Storing Java objects as column data in a table" on page 76.</p>
VERSION_3	<ul style="list-style-type: none"> The default value of the LANGUAGE connection property is us_english. If the CHARSET connection property does not specify a character set, jConnect uses the database's default character set. The default value for CHARSET_CONVERTER is the PureConverter class. By default, Statement.cancel() cancels the object it is invoked on and any other Statement objects that have begun to execute and are waiting for results. 	<p>See the comments for VERSION_2.</p>
VERSION_2	<ul style="list-style-type: none"> The default value of the LANGUAGE connection property is us_english. If the CHARSET connection property does not specify a character set, the default character set is iso_1. The default value for CHARSET_CONVERTER is the TruncationConverter class, unless the CHARSET connection property specifies a multibyte or 8-bit character set, in which case the default CHARSET_CONVERTER is the PureConverter class. By default, Statement.cancel() cancels the object it is invoked on and any other Statement objects that have begun to execute and are waiting for results. 	<p>The LANGUAGE connection property determines the language in which messages from jConnect and the server appear.</p> <p>For information on the CHARSET and CHARSET_CONVERTER connection classes, see "jConnect character-set converters" on page 34.</p> <p>The VERSION_2 default behavior of Statement.cancel() is not JDBC-compliant. Use CANCEL_ALL to set the behavior of Statement.cancel(). See "CANCEL_ALL connection property" on page 10.</p>

The version values are constant values from the `SybDriver` class. When referring to the version constant, use this syntax:

```
com.sybase.jdbcx.SybDriver.VERSION_5
```

Use `SybDriver.setVersion()` to set the `jConnect` version. The `setVersion` method will affect the `jConnect` default behavior for all `Connections` created by the `SybDriver` object. However, you can use the `JCONNECT_VERSION` connection property to set version-specific behavior for individual connections. The following code samples show how to load the `jConnect` driver and set the version:

```
import com.sybase.jdbcx.SybDriver;
SybDriver sybDriver = (SybDriver)
    Class.forName
        ("com.sybase.jdbc2.jdbc.SybDriver").newInstance();
sybDriver.setVersion
    (com.sybase.jdbcx.SybDriver.VERSION_5);
DriverManager.registerDriver(sybDriver);
```

You can call `setVersion()` multiple times to change the version setting. New connections inherit the behavior associated with the version setting at the time the connection was made. Changing the version setting during a session does not affect the current connection. `jConnect` provides a `com.sybase.jdbcx.SybDriver.VERSION_LATEST` constant which can be used to ensure that you are always requesting the highest version value possible for the `jConnect` driver you're using. However, by setting the version to `VERSION_LATEST`, you may see behavior changes, if you replace your `jConnect` driver with a newer one.

As described in the next section, you can use `JCONNECT_VERSION` to override the `SybDriver` version setting and specify a different version setting for a specific connection.

For wide table, unichar, and date/time datatype support with `jConnect 5.5`:

To enable wide table, unichar, and date/time datatype support with Adaptive Server version 12.5 and later, (date/time datatype support is only available in Adaptive Server version 12.5.1 and later), use `SybDriver.setVersion()` to set the `jConnect` version to 6:

```
sybDriver.setVersion
    (com.sybase.jdbcx.SybDriver.VERSION_6);
```

JCONNECT_VERSION connection property

Use JCONNECT_VERSION to specify the version setting for a specific session. You can set JCONNECT_VERSION to an integer value of “2,” “3,” “4,” “5,” or “6” depending on the characteristics you want (see Table 2-1).

CANCEL_ALL connection property

CANCEL_ALL is a Boolean-valued connection property for specifying the behavior of the `Statement.cancel()` method.

Note In jConnect version 4.0 and earlier, the default for CANCEL_ALL is *true*. In jConnect version 4.1 and later, to comply with the JDBC specification, if you set the connection property JCONNECT_VERSION to “4” or above, the default setting for CANCEL_ALL is *false*.

The settings for CANCEL_ALL have the following effect on `Statement.cancel()`:

- If CANCEL_ALL is *false*, invoking `Statement.cancel()` cancels only the `Statement` object it is invoked on. Thus, if `stmtA` is a `Statement` object, `stmtA.cancel()` cancels the execution of the SQL statement contained in `stmtA` in the database, but no other statements are affected. `stmtA` is canceled whether it is in cache waiting to execute or has started to execute and is waiting for results.
- If CANCEL_ALL is *true*, invoking `Statement.cancel()` cancels not only the object it is invoked on, but also any other `Statement` objects on the same connection that have executed and are waiting for results.

The following example sets CANCEL_ALL to *false*. In the example, *props* is a `Properties` object for specifying connection properties.

```
props.put("CANCEL_ALL", "false");
```

Note To cancel the execution of all `Statement` objects on a connection, regardless of whether or not they have begun execution on the server, use the extension method `SybConnection.cancel()`.

Invoking the jConnect driver

To register and invoke the Sybase jConnect driver, use either of the following two suggested methods:

- 1 Use a `Class.forName` call as in the following example:

```
Class.forName("com.sybase.jdbc2.jdbc.SybDriver")
        .newInstance();
```

- 2 Add the jConnect driver to the `jdbc.drivers` system property. At initialization, the `DriverManager` class attempts to load the drivers listed in `jdbc.drivers`. This is less efficient than the `Class.forName` call approach. You can list multiple drivers in this property, separated with a colon (:). The following code samples show how to add a driver to `jdbc.drivers` within a program:

```
Properties sysProps = System.getProperties();
String drivers = "com.sybase.jdbc2.jdbc.SybDriver";
String oldDrivers =
sysProps.getProperty("jdbc.drivers");
if (oldDrivers != null)
    drivers += ":" + oldDrivers;
sysProps.put("jdbc.drivers", drivers.toString());
```

Note `System.getProperties()` is not allowed for Java applets. Use the `Class.forName()` method, instead.

Establishing a connection

This section describes how to establish a connection to an Adaptive Server or Adaptive Server Anywhere database using jConnect.

Setting connection properties

Table 2-2 lists the connection properties for jConnect and indicates their default values. *You must set the connection properties before you make a connection.*

There are two ways to set the driver connection properties:

- Use the `DriverManager.getConnection()` method in your application.
- When you define the URL.

Note Driver connection properties set in the URL do not override any corresponding connection properties set in the application using the `DriverManager.getConnection()` method.

To obtain a current list of properties for any driver, use the `Driver.getDriverPropertyInfo(String url, Properties props)`, which returns an array of `DriverPropertyInfo` objects. The array lists:

- Driver properties
- Current settings on which the driver properties are based
- The URL and props passed in

Driver connection property names are not case-sensitive (jConnect uses the `String.equalsIgnoreCase(String)` method to compare property names).

Table 2-2: Connection properties

Property	Description	Default value
APPLICATIONNAME	A user-defined property. The server side can be programmed to interpret the value given to this property.	Null
BE_AS_JDBC_COMPLIANT_ AS_POSSIBLE (only for jConnect version 5.5 or later)	Adjusts other properties to ensure that jConnect methods respond in a way that is as compliant as possible with the JDBC 2.1 standard. These properties are affected (and overridden) when this property is set to <i>true</i> : <ul style="list-style-type: none">• CANCEL_ALL (set to <i>false</i>)• LANGUAGE_CURSOR (set to <i>true</i> for jConnect 4.x, set to <i>false</i> for jConnect 5.x)• SELECT_OPENS_CURSOR (set to <i>true</i>)• FAKE_METADATA (set to <i>true</i>)• GET_BY_NAME_USES_COLUMN_LABEL (set to <i>false</i>)	False
CANCEL_ALL	Determines the behavior of the <code>Statement.cancel()</code> method. See “CANCEL_ALL connection property” on page 10.	Depends on version setting (see “Setting the jConnect version” on page 6.

Property	Description	Default value
CHARSET	Specifies the character set for strings passed to the database. If the CHARSET value is null, jConnect will use the server's default character set to send string data to the server. If you specify a CHARSET, the database must be able to handle characters in that format. If the database cannot do so, a message will be generated indicating that character conversion cannot be properly completed. Note that when using jConnect 5.5 with unichar enabled, jConnect will detect when a client is trying to send characters to the server that cannot be represented in the character set that is being used for the Connection. When that occurs, jConnect will send the character data to the server as unichar data; this allows clients to insert unicode data into unichar/univarchar columns and parameters.	Null
CHARSET_CONVERTER_CLASS	Use this property to specify the character-set converter class you want jConnect to use. jConnect uses the version setting from SybDriver.setVersion(), or the version passed in with the JCONNECT_VERSION property, to determine the default character-set converter class to use. See "Selecting a character-set converter" on page 35 for details.	Version dependent
CLASS_LOADER	Set this property to a DynamicClassLoader object that you create. The DynamicClassLoader is used to load Java classes that are stored in the database, but which are not in the CLASSPATH at application start-up time. See "Dynamic class loading" on page 80 for more information.	Null
CONNECTION_FAILOVER	For use with the Java Naming and Directory Interface (JNDI). See "CONNECTION_FAILOVER connection property" on page 25.	True
DISABLE_UNICHAR_SENDING (Only for jConnect version 5.5 or later)	When a client application sends unichar characters to the server (along with non-unichar characters), there is a slight performance hit for any character data sent to the database. Because this property defaults to <i>true</i> , clients who wish to send unichar data to the database <i>must</i> set this property value to <i>false</i> . See "Using jConnect to pass Unicode data" on page 33.	True

Property	Description	Default value
DISABLE_UNPROCESSED_PARAM_WARNINGS	During results processing for a stored procedure, jConnect often reads return values other than row data. If you do not process the return value, jConnect raises a warning. To disable these warnings (which can help performance), set this property to <i>true</i> .	False
DYNAMIC_PREPARE	Determines whether dynamic SQL prepared statements are precompiled in the database. See “DYNAMIC_PREPARE connection property” on page 118.	False
ESCAPE_PROCESSING_DEFAULT	By default, jConnect parses all SQL statements submitted to the database for valid JDBC function escapes. If your application is not going to use JDBC function escapes in its SQL calls, you can set this connection property to <i>false</i> to circumvent this processing. This can provide a slight performance benefit.	True
EXPIRESTRING	A read-only property that contains the license expiration date. Expiration is never except for evaluation copies of jConnect.	Never
FAKE_METADATA (Only for jConnect version 5.5 or later)	<p>When you call the <code>ResultSetMetaData</code> methods <code>getCatalogName</code>, <code>getSchemaName</code>, and <code>getTableName</code> and this property is set to <i>true</i>, the call will return empty strings (“”) because the server does not supply useful metadata.</p> <p>When this property is set to <i>false</i>, calling these methods throws a “Not Implemented” <code>SQLException</code>.</p> <hr/> <p>Note If you have enabled wide tables and are using an Adaptive Server 12.5 or later, this property setting is ignored, because the server does supply useful metadata.</p> <hr/>	False

Property	Description	Default value
GET_BY_NAME_USES_COLUMN_LABEL (Only for jConnect version 5.5 or later)	<p>Provides backward compatibility with versions of jConnect earlier than 5.5.</p> <p>With Adaptive Server version 12.5, jConnect has access to more metadata than was previously available. Previous to version 12.5, column name and column alias meant the same thing. jConnect can now differentiate between the two when used with a 12.5 or later Adaptive Server with wide tables enabled.</p> <p>To preserve backward compatibility, set this property to <i>true</i>.</p> <p>If you want calls to <code>getBytes</code>, <code>getInt</code>, <code>get*</code> (String <code>columnName</code>) to look at the actual name for the column (called for in the JDBC 2.0 specification), set this property to <i>false</i>.</p>	True
HOSTNAME	The name of the current host.	None The max length is 30 characters and if exceeded, it will be truncated to 30
HOSTPROC	Identifies the application's process on the host machine.	None
IGNORE_DONE_IN_PROC	When set to <i>true</i> , intermediate update results (as in stored procedures) are not returned, only the final result set.	False
IS_CLOSED_TEST	Allows you to specify what query, if any, is sent to the database when <code>Connection.isClosed()</code> is called. For additional information, see the "Connection.isClosed() and the IS_CLOSED_TEST connection property" on page 95.	Null
JCONNECT_VERSION	Use this property to set version-specific characteristics. See "JCONNECT_VERSION connection property" on page 10.	5
LANGUAGE	Set this property for error messages returned from the server and for jConnect messages. It must match a language in <i>syslanguages</i> .	Version dependent (see "Setting the jConnect version" on page 6).
LANGUAGE_CURSOR	<p>Set this property to <i>true</i> if you want jConnect to use "language cursors" instead of "protocol cursors."</p> <p>See "Cursor performance" on page 120.</p>	False

Property	Description	Default value
LITERAL_PARAMS	<p>Use this property primarily when using jConnect with Adaptive Server Anywhere version 5.5. ASA 5.5 requires you to send prepared statement parameters as literals. For all other Sybase databases (including ASA 6 and later), this property can be set to <i>false</i>.</p> <p>When set to <i>true</i>, any parameters set by the <code>setXXX</code> methods in the <code>PreparedStatement</code> interface are inserted literally into the SQL statement when it is executed.</p> <p>If set to <i>false</i>, parameter markers are left in the SQL statement and the parameter values are sent to the server separately.</p>	False
USE_METADATA	<p>When set to <i>true</i>, a <code>DatabaseMetaData</code> object will be created and initialized when you establish a connection. The <code>DatabaseMetaData</code> object is necessary to connect to a specified database.</p> <p>jConnect uses <code>DatabaseMetaData</code> for some features, including Distributed Transaction Management support (JTA/JTS) and dynamic class loading (DCL).</p> <p>If you receive error 010SJ, which indicates that your application requires metadata, install the stored procedures for returning metadata that come with jConnect (see “Installing Stored Procedures” in Chapter 3 of the <i>jConnect for JDBC Installation Guide</i>).</p>	true
PACKETSIZE	Network packet size.	512
PASSWORD	<p>Login password.</p> <p>Set automatically if using the <code>getConnection(String, String, String)</code> method, or explicitly if using <code>getConnection(String, Props)</code>.</p>	None
PRELOAD_JARS	<p>A comma-separated list of jar file names that are associated with the <code>CLASS_LOADER</code> that you specify. These jars are loaded at connect time, and are available for use by any other Connection using the same jConnect driver. See “Preloading JARS” on page 83 for more information.</p>	Null

Property	Description	Default value
PROTOCOL_CAPTURE	The PROTOCOL_CAPTURE connection property is used to specify a file for capturing TDS communication between an application and an Adaptive Server.	Null
PROXY	Gateway address. For the HTTP protocol, the URL is <i>http://host:port</i> . To use the HTTPS protocol that supports encryption, the URL is <i>https://host:port/servlet_alias</i> .	None
REMOTEPWD	Remote server passwords for access through server-to-server remote procedure calls. See “Performing server-to-server remote procedure calls” on page 45.	None
REPEAT_READ	Determines whether the driver keeps copies of columns and output parameters so that columns can be read out of order or repeatedly. See “REPEAT_READ connection property” on page 112.	True
REQUEST_HA_SESSION	This property indicates whether the connecting client wants to begin an HA failover session with a version 12 or later Adaptive Server configured for failover. See “Implementing high availability failover support” on page 40. Setting this property to <i>true</i> causes jConnect to attempt a failover login. If you do not set this connection property, a failover session will not start, even if the server is configured for failover. You cannot reset the property once a connection has been made. If you want more flexibility for requesting failover sessions, code the client application to set REQUEST_HA_SESSION at runtime.	False
RMNAME	This property is used to set the Resource Manager name when using distributed transactions (XA). This property will override a Resource Manager name that may be set in an LDAP server entry. See “Distributed transaction management support” on page 91 for more information.	Null

Property	Description	Default value
SECONDARY_SERVER_HOSTPORT	Sets the hostname and port for the secondary server when the client is using an HA failover session. The value for this property should be in the form of <code>hostName:portNumber</code> . This property is ignored unless you have also set <code>REQUEST_HA_SESSION</code> to <code>true</code> . See “Implementing high availability failover support” on page 40 for more information.	Null
SELECT_OPENS_CURSOR	<p>If set to <code>true</code>, calls to <code>Statement.executeQuery()</code> will automatically generate a cursor when the query contains a “FOR UPDATE” clause.</p> <p>If you have previously called <code>Statement.setFetchSize()</code> or <code>Statement.setCursorName()</code> on the same statement, a setting of <code>true</code> for <code>SELECT_OPENS_CURSOR</code> has no effect.</p> <hr/> <p>Note You may experience some performance degradation when <code>SELECT_OPENS_CURSOR</code> is set to <code>true</code>.</p> <hr/> <p>See “Using cursors with result sets” on page 48 for more information on using cursors with <code>jConnect</code>.</p>	False
SERIALIZE_REQUESTS	If set to <code>true</code> , <code>jConnect</code> waits for responses from the server before sending additional requests.	False
SERVICENAME	The name of a back-end database server that a <code>DirectConnect</code> gateway serves. Also used to indicate which database should be used upon connecting to Adaptive Server Anywhere.	None
SESSION_ID	When this property is set, <code>jConnect</code> assumes that an application is trying to resume communication on an existing TDS session held open by the TDS-tunnelling gateway. <code>jConnect</code> skips the login negotiations and forwards all requests from the application to the specified session ID.	Null

Property	Description	Default value
SESSION_TIMEOUT	Use this property to specify the amount of time (in seconds) that an http-tunnelled session (created using the jConnect TDS-tunnelling servlet) will remain alive while idle. After the specified time, the connection will be automatically closed. For more information about the TDS-tunnelling servlet, see page 134.	Null
SQLINITSTRING	Use this property to define a set of commands to be passed to the database server when a Connection is opened. These must be SQL commands that can be executed using the <code>Statement.executeUpdate()</code> method.	Null
STREAM_CACHE_SIZE	Maximum size used to cache statement response streams.	Null (unlimited cache size)
SYB SOCKET_FACTORY	Use this property to enable jConnect to use your custom socket implementation. Set SYB SOCKET_FACTORY either to: <ul style="list-style-type: none"> • The name of a class that implements <code>com.sybase.jdbcx.SybSocketFactory</code>; or • “DEFAULT,” which instantiates a new <code>java.net.Socket()</code> See “Implementing custom socket plug-ins” on page 28.	Null
USER	Login ID. Set automatically if using the <code>getConnection(String, String, String)</code> method, or explicitly if using <code>getConnection(String, Props)</code> .	None
VERSIONSTRING	Read-only version information for the JDBC driver.	jConnect driver version

The following code is an example of setting connection properties. The sample programs provided with jConnect also contain examples of setting these properties.

```
Properties props = new Properties();
props.put("user", "userid");
props.put("password", "user_password");
/*
 * If the program is an applet that wants to access
 * a server that is not on the same host as the
 * web server, then it uses a proxy gateway.
 */
```

```
props.put("proxy", "localhost:port");
/*
 * Make sure you set connection properties before
 * attempting to make a connection. You can also
 * set the properties in the URL.
 */
Connection con = DriverManager.getConnection
    ("jdbc:sybase:Tds:host:port", props);
```

Connecting to Adaptive Server

In your Java application, define a URL using the jConnect driver to connect to an Adaptive Server. The basic format of the URL is:

```
jdbc:sybase:Tds:host:port
```

where:

`jdbc:sybase` identifies the driver.

`Tds` is the Sybase communication protocol for Adaptive Server.

`host:port` is the Adaptive Server host name and listening port. See *\$SYBASE/interfaces* (UNIX) or *%SYBASE%\ini\sql.ini* (Windows) for the entry that your database or Open Server application uses. Obtain the *host:port* from the “query” entry.

You can connect to a specific database using this format:

```
jdbc:sybase:Tds:host:port/database
```

Note To connect to a specific database using Adaptive Server Anywhere 6.x and later or DirectConnect, use the `SERVICENAME` connection property to specify the database name instead of “/database.”

Example

The following code creates a connection to an Adaptive Server on host “myserver” listening on port 3697:

```
SysProps.put("user", "userid");
SysProps.put("password", "user_password");
String url = "jdbc:sybase:Tds:myserver:3697";
Connection con =
    DriverManager.getConnection(url, SysProps);
```

URL connection property parameters

You can specify the values for the jConnect driver connection properties when you define a URL.

Note Driver connection properties set in the URL do not override any corresponding connection properties set in the application using the `DriverManager.getConnection()` method.

To set a connection property in the URL, append the property name and its value to the URL definition. Use this syntax:

```
jdbc:sybase:Tds:host:port/database?  
property_name=value
```

To set multiple connection properties, append each additional connection property and value, preceded by “&.” For example:

```
jdbc:sybase:Tds:myserver:1234/mydatabase?  
LITERAL_PARAMS=true&PACKETSIZE=512&HOSTNAME=myhost
```

If the value for one of the connection properties contains “&,” precede the “&” in the connection property value with a backslash (\). For example, if your host name is “a&bhost,” use this syntax:

```
jdbc:sybase:Tds:myserver:1234/mydatabase?  
LITERAL_PARAMS=true&PACKETSIZE=512&HOSTNAME=  
a\&bhost
```

Do not use quotes for connection property values, even if they are strings. For example, use:

```
HOSTNAME=myhost
```

not:

```
HOSTNAME="myhost "
```

Connecting to Adaptive Server Anywhere

To use jConnect with Adaptive Server Anywhere, you should upgrade to Adaptive Server Anywhere version 6.x and later.

Connecting to Adaptive Server Anywhere 5.x.x

If you have to connect to Adaptive Server Anywhere version 5.x.x through jConnect, you must run the Adaptive Server Anywhere Open Server Gateway dbos50, which is distributed with Adaptive Server Anywhere.

- 1 Install Open Server Gateway 5.5.x3 or later and the Open Server DLLs. Use Open Server DLLs, version 11.1.
- 2 Add an entry for the gateway to your `%SYBASE%\ini\sql.ini` file (using, for example, sqledit).
- 3 Start the gateway by entering:

```
start dbos50 gateway-demo
```

where *gateway-demo* is the gateway name defined in step 2.

- 4 When the Open Server Gateway is running, you can define a connection as follows:

```
jdbc:sybase:Tds:host:port
```

host is the host name where the Adaptive Server Anywhere and Open Server gateway is running, and *port* is the port number defined in *sql.ini*.

Note To support multiple Adaptive Server Anywhere databases, use sqledit to add an entry with a different port for each database, then run the Open Server Gateway for each database.

Connecting to a server using JNDI

In jConnect 4.0 and later, you can use the Java Naming and Directory Interface (JNDI) to provide connection information, which offers:

- A centralized location where you can specify host names and ports for connecting to a server. You do not need to hard code a specific host and port number in an application.
- A centralized location where you can specify connection properties and a default database for all applications to use.
- The jConnect `CONNECTION_FAILOVER` property for handling unsuccessful connection attempts. When `CONNECTION_FAILOVER` is set to true, jConnect attempts to connect to a sequence of host/port server addresses in the JNDI name space until one succeeds.

To use jConnect with JNDI, you need to make sure that certain information is available in any directory service that JNDI accesses and that required information is set in the `javax.naming.Context` class. This section covers the following topics:

- Connection URL for using JNDI
- Required directory service information
- `CONNECTION_FAILOVER` connection property
- Providing JNDI context information

Connection URL for using JNDI

To specify that jConnect use JNDI to obtain connection information, place “jndi” as the URL’s subprotocol after “sybase”:

```
jdbc:sybase:jndi:protocol-information-for-use-with-JNDI
```

Anything that follows “jndi” in the URL is handled through JNDI. For example, to use JNDI with the Lightweight Directory Access Protocol (LDAP), you might enter:

```
jdbc:sybase:jndi:ldap://LDAP_hostname:port_number/servername=
Sybase11,o=MyCompany,c=US
```

This URL tells JNDI to obtain information from an LDAP server, gives the host name and port number of the LDAP server to use, and provides the name of a database server in an LDAP-specific form.

Required directory service information

When you use JNDI with jConnect, JNDI needs to return the following information for the target database server:

- A host name and port number to connect to
- The name of the database to use
- Any connection properties that individual applications are not allowed to set on their own

This information needs to be stored according to a fixed format in any directory service used for providing connection information. The required format consists of a numerical object identifier (OID), which identifies the type of information being provided (for example, the destination database), followed by the formatted information (see “Example 1” on page 23).

Note You can use the alias name for to reference the attribute instead of the OID. See “Example 2” on page 24.

Table 2-3 shows the required formatting.

Table 2-3: Directory service information for JNDI

Attribute description	Alias	OID (object_id)
Interfaces entry replacement in LDAP directory services	sybaseServer	1.3.6.1.4.1.897.4.1.1
Collection point for sybaseServer LDAP attributes	sybaseServer	1.3.6.1.4.1.897.4.2
Version Attribute	sybaseVersion	1.3.6.1.4.1.897.4.2.1
Servename Attribute	sybaseServer	1.3.6.1.4.1.897.4.2.2
Service Attribute	sybaseService	1.3.6.1.4.1.897.4.2.3
Status Attribute	sybaseStatus	1.3.6.1.4.1.897.4.2.4
<i>Address Attribute</i>	sybaseAddress	1.3.6.1.4.1.897.4.2.5
Security Mechanism Attribute	sybaseSecurity	1.3.6.1.4.1.897.4.2.6
Retry Count Attribute	sybaseRetryCount	1.3.6.1.4.1.897.4.2.7
Loop Delay Attribute	sybaseRetryDelay	1.3.6.1.4.1.897.4.2.8
<i>jConnect Connection Protocol</i>	sybaseJconnectProtocol	1.3.6.1.4.1.897.4.2.9
<i>jConnect Connection Property</i>	sybaseJconnectProperty	1.3.6.1.4.1.897.4.2.10
<i>Database Name</i>	sybaseDatabasename	1.3.6.1.4.1.897.4.2.11
High Availability Failover Servename Attribute	sybaseHAServename	1.3.6.1.4.1.897.4.2.15
ResourceManager Name	sybaseResourceManager Name	1.3.6.1.4.1.897.4.2.16
ResourceManager Type	sybaseResourceManager Type	1.3.6.1.4.1.897.4.2.17
JDBCDataSource Interface	sybaseJdbcDataSource-Interface	1.3.6.1.4.1.897.4.2.18
ServerType	sybaseServerType	1.3.6.1.4.1.897.4.2.19

Note Attributes in italics are required.

The following examples show connection information entered for the database server SYBASE11 under an LDAP directory service. Example 1 uses the attribute's OID, Example 2 uses the attribute's alias, which is case insensitive. You can use either the OID or the alias.

Example 1

```
dn: servername=SYBASE11,o=MyCompany,c=US
servername:SYBASE11
1.3.6.1.4.1.897.4.2.5:TCP#1#giotto 1266
1.3.6.1.4.1.897.4.2.5:TCP#1#giotto 1337
1.3.6.1.4.1.897.4.2.5:TCP#1#standby1 4444
1.3.6.1.4.1.897.4.2.10:REPEAT_READ=false&PACKETSIZE=1024
1.3.6.1.4.1.897.4.2.10:CONNECTION_FAILOVER=true
1.3.6.1.4.1.897.4.2.11:pubs2
1.3.6.1.4.1.897.4.2.9:Tds
```

Example 2

```
dn: servername=SYBASE11,o=MyCompany,c=US
servername:SYBASE11
sybaseAddress:TCP#1#giotto 1266
sybaseAddress:TCP#1#giotto 1337
sybaseAddress:TCP#1#standby1 4444
sybaseJconnectProperty:REPEAT_READ=false&PACKETSIZE=1024
sybaseJconnectProperty:CONNECTION_FAILOVER=true
sybaseDatabasename:pubs2
sybaseJconnectProtocol:Tds
```

In these examples, SYBASE11 can be accessed through either port 1266 or port 1337 on host “giotto” and it can be accessed through port 4444 on host “standby1.” Two connection properties, REPEAT_READ and PACKETSIZE, are set within one entry. The CONNECTION_FAILOVER connection property is set as a separate entry. Applications connecting to SYBASE11 are initially connected with the pubs2 database. You do not need to specify a connection protocol, but if you do, you must enter the attribute as “Tds”, not “TDS”.

CONNECTION_FAILOVER connection property

CONNECTION_FAILOVER is a Boolean-valued connection property you can use when jConnect uses JNDI to get connection information.

If `CONNECTION_FAILOVER` is set to true, jConnect makes multiple attempts to connect to a server. If one attempt to connect to a host and port number associated with a server fails, jConnect uses JNDI to get the next host and port number associated with the server and attempts to connect through them. Connection attempts proceed sequentially through all the hosts and ports associated with a server.

For example, suppose `CONNECTION_FAILOVER` is set to true, and a database server is associated with the following hosts and port numbers, as in the earlier LDAP example:

```
1.3.6.1.4.1.897.4.2.5:TCP#1#giotto 1266
1.3.6.1.4.1.897.4.2.5:TCP#1#giotto 1337
1.3.6.1.4.1.897.4.2.5:TCP#1#standby 4444
```

To get a connection to the server, jConnect tries to connect to the host “giotto” at port 1266. If this fails, jConnect tries port 1337 on “giotto.” If this fails, jConnect tries to connect to host “standby1” through port 4444.

The default for `CONNECTION_FAILOVER` is true.

If `CONNECTION_FAILOVER` is set to false, jConnect attempts to connect to an initial host and port number. If the attempt fails, jConnect throws a SQL exception and does not try again.

Providing JNDI context information

To use jConnect with JNDI, a developer should be familiar with the JNDI specification from Sun Microsystems, available from the Web:

<http://java.sun.com/products/jndi>

In particular, the developer needs to make sure that required initialization properties are set in `javax.naming.directory.DirContext` when JNDI and jConnect are used together. These properties can be set either at the system level or at runtime.

Two key properties are:

- `Context.INITIAL_CONTEXT_FACTORY`

This property takes the fully qualified class name of the initial context factory for JNDI to use. This determines the JNDI driver that is used with the URL specified in the `Context.PROVIDER_URL` property.

- `Context.PROVIDER_URL`

This property takes the URL of the directory service that the driver (for example, the LDAP driver) is to access. The URL should be a string, such as “ldap://ldaphost:427”.

The following example shows how to set context properties at runtime and how to get a connection using JNDI and LDAP. In the example, the `INITIAL_CONTEXT_FACTORY` context property is set to invoke Sun Microsystems’s implementation of an LDAP service provider. The `PROVIDER_URL` context property is set to the URL of an LDAP directory service located on the host “ldap_server1” at port 983.

```
Properties props = new Properties();

/* We want to use LDAP, so INITIAL_CONTEXT_FACTORY is set to the
 * class name of an LDAP context factory. In this case, the
 * context factory is provided by Sun’s implementation of a
 * driver for LDAP directory service.
 */
props.put(Context.INITIAL_CONTEXT_FACTORY,
    "com.sun.jndi.ldap.LdapCtxFactory");

/* Now, we set PROVIDER_URL to the URL of the LDAP server that
 * is to provide directory information for the connection.
 */
props.put(Context.PROVIDER_URL, "ldap://ldap_server1:983");

/* Set up additional context properties, as needed. */
props.put("user", "xyz");
props.put("password", "123");

/* get the connection */
Connection con = DriverManager.getConnection
    ("jdbc:sybase:jndi:ldap://ldap_server1:983" +
    "/servername=Sybase11,o=MyCompany,c=US",props);
```

The connection string passed to `getConnection()` contains LDAP-specific information, which the developer must provide.

When JNDI properties are set at runtime, as in the preceding example, `jConnect` passes them to JNDI to be used in initializing a server, as in the following `jConnect` code:

```
javax.naming.directory.DirContext ctx =
    new javax.naming.directory.InitialDirContext(props);
```

`jConnect` then obtains the connection information it needs from JNDI by invoking `DirContext.getAttributes()`, as in the following example, where `ctx` is a `DirContext` object:

```
javax.naming.directory.Attributes attrs =
    ctx.getAttributes(ldap://ldap_server1:983/servername=
        Sybase11, SYBASE_SERVER_ATTRIBUTES);
```

In the example, SYBASE_SERVER_ATTRIBUTES is an array of strings defined within jConnect. The array values are the OIDs for the required directory information listed in Table 2-3.

Implementing custom socket plug-ins

This section discusses how to plug a custom socket implementation into an application to customize the communication between a client and server. `javax.net.ssl.SSLSocket` is an example of a socket that you could customize to enable encryption.

`com.sybase.jdbcx.SybSocketFactory` is a Sybase extension interface that contains the `createSocket(String, int, Properties)` method that returns a `java.net.Socket`. For a jConnect version 4.1 or later driver to load a custom socket, an application must:

- Implement this interface
- Define the `createSocket(..)` method

jConnect uses the new socket for its subsequent input/output operations. Classes that implement `SybSocketFactory` create sockets and provide a general framework for the addition of public socket-level functionality.

```
/**
 * Returns a socket connected to a ServerSocket on the named host,
 * at the given port.
 * @param host    the server host
 * @param port    the server port
 * @param props   Properties passed in through the connection
 * @returns Socket
 * @exception IOException, UnknownHostException
 */
public java.net.Socket createSocket(String host, int port, Properties props)
throws IOException, UnknownHostException;
```

Passing in properties allows instances of `SybSocketFactory` to use connection properties to implement an intelligent socket.

When you implement `SybSocketFactory` to produce a socket, the same application code can use different kinds of sockets by passing the different kinds of factories or pseudo-factories that create sockets to the application. You can customize factories with parameters used in socket construction. For example, you could customize factories to return sockets with different networking time outs or security parameters already configured. The sockets returned to the application can be subclasses of `java.net.Socket` to directly expose new APIs for features such as compression, security, record marking, statistics collection, or firewall tunnelling (`javax.net.SocketFactory`).

Note `SybSocketFactory` is intended to be an overly simplified `javax.net.SocketFactory`, enabling applications to bridge from `java.net.*` to `javax.net.*` if desired.

To use a custom socket with `jConnect`:

- 1 Provide a Java class that implements `com.sybase.jdbcx.SybSocketFactory`. See “Creating and configuring a custom socket” on page 29.
- 2 Set the `SYB_SOCKET_FACTORY` connection property so that `jConnect` can use your implementation to obtain a socket.

SYB_SOCKET_FACTORY connection property

To use a custom socket with `jConnect`, set the `SYB_SOCKET_FACTORY` connection property to a string that is either:

- The name of a class that implements `com.sybase.jdbcx.SybSocketFactory`
or
- `DEFAULT`, which instantiates a new `java.net.Socket()`

See “Setting connection properties” on page 11 for instructions on how to set `SYB_SOCKET_FACTORY`.

Creating and configuring a custom socket

Once `jConnect` obtains a custom socket, it uses the socket to connect to a server. Any configuration of the socket must be completed before `jConnect` obtains it.

This section explains how to plug in an SSL socket implementation, such as `javax.net.ssl.SSLSocket`, with `jConnect`.

Note Currently, only Adaptive Server version 12.5 and later supports SSL.

The following example shows how an implementation of SSL can create an instance of `SSLSocket`, configure it, and then return it. In the example, the `MySSLSocketFactory` class implements `SybSocketFactory` and extends `javax.net.ssl.SSLSocketFactory` to implement SSL. It contains two `createSocket` methods—one for `SSLSocketFactory` and one for `SybSocketFactory`—that:

- Create an SSL socket
- Invoke `SSLSocket.setEnabledCipherSuites()` to specify the cipher suites available for encryption
- Return the socket to be used by `jConnect`

Example

```
public class MySSLSocketFactory extends SSLSocketFactory
    implements SybSocketFactory
{
    /**
     * Create a socket, set the cipher suites it can use, return
     * the socket.
     * Demonstrates how either suites could be hard-coded into the
     * implementation.
     *
     * See javax.net.SSLSocketFactory#createSocket
     */
    public Socket createSocket(String host, int port)
        throws IOException, UnknownHostException
    {
        // Prepare an array containing the cipher suites that are to
        // be enabled.
        String enableThese[] =
        {
            "SSL_DH_DSS_EXPORT_WITH_DES40_CBC_SHA",
            "SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5",
            "SSL_DH_RSA_EXPORT_WITH_DES40_CBC_SHA"
        }
        ;

        Socket s =
            SSLSocketFactory.getDefault().createSocket(host, port);
```



```

        ((SSLSocket)s).setEnabledCipherSuites(enableThese);
    }
    return s;
}
/**
 * Return an SSLSocket.
 * Demonstrates how to set cipher suites based on connection
 * properties like:
 * Properties _props = new Properties();
 * Set other url, password, etc. properties.
 * _props.put("CIPHER_SUITES_1",
 *     "SSL_DH_DSS_EXPORT_WITH_DES40_CBC_SHA");
 * _props.put("CIPHER_SUITES_2",
 *     "SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5");
 * _props.put("CIPHER_SUITES_3",
 *     "SSL_DH_RSA_EXPORT_WITH_DES40_CBC_SHA");
 * _conn = _driver.getConnection(url, _props);
 *
 * See com.sybase.jdbcx.SybSocketFactory#createSocket
 */
public Socket createSocket(String host, int port,
    Properties props)
    throws IOException, UnknownHostException
{
    // check to see if cipher suites are set in the connection
    // properties
    Vector cipherSuites = new Vector();
    String cipherSuiteVal = null;
    int cipherIndex = 1;
    do
    {
        if((cipherSuiteVal = props.getProperty("CIPHER_SUITES_"
            + cipherIndex++)) == null)
        {
            if(cipherIndex <= 2)
            {
                // No cipher suites available
                // return what the object considers its default
                // SSLSocket, with cipher suites enabled.
                return createSocket(host, port);
            }
            else
            {
                // we have at least one cipher suite to enable
                // per request on the connection
                break;
            }
        }
    }
}

```

```
        else
        {
            // add to the cipher suit Vector, so that
            // we may enable them together
            cipherSuites.addElement(cipherSuiteVal);
        }
    }
    while(true);
    // lets you create a String[] out of the created vector
    String enableThese[] = new String[cipherSuites.size()];
    cipherSuites.copyInto(enableThese);
    // enable the cipher suites
    Socket s =
        SSLSocketFactory.getDefault().createSocket
            (host, port);
    ((SSLSocket)s).setEnabledCipherSuites(enableThese);
    // return the SSLSocket
    return s;
}
// other methods
}
```

Since `jConnect` requires no information about the kind of socket it is, you must complete any configuration before you return a socket.

For additional information, see:

- *Encrypt.java* – located in the *sample2* (`jConnect 5.x`) subdirectories of your `jConnect` directory, this sample shows you how to use the `SybSocketFactory` interface with `jConnect` applications.
- *MySSLSocketFactory.java* – also located in the *sample2* (`jConnect 5.x`) subdirectories of your `jConnect` directory, this is a sample implementation of the `SybSocketFactory` interface that you can plug in to your application and use.

Handling internationalization and localization

This section discusses internationalization and localization issues relevant to `jConnect`.

Using jConnect to pass Unicode data

In Adaptive Server version 12.5 and later, database clients can take advantage of the `unichar` and `univarchar` datatypes. The two datatypes allow for the efficient storage and retrieval of Unicode data.

Quoting from the Unicode Standard, version 2.0:

“The Unicode Standard is a fixed-width, uniform encoding scheme for encoding characters and text. The repertoire of this international character code for information processing includes characters for the major scripts of the world, as well as technical symbols in common. The Unicode character encoding treats alphabetic characters, ideographic characters, and symbols identically, which means they can be used in any mixture and with equal facility. The Unicode Standard is modeled on the ASCII character set, but uses a 16-bit encoding to support full multilingual text.”

This means that the user can designate database table columns to store Unicode data, regardless of the default character set of the server.

Note In Adaptive Server version 12.5 through 12.5.0.3, the server had to have a default character set of `utf-8` in order to use the unicode datatypes. However, in Adaptive Server 12.5.1 and later, database users can use `unichar` and `univarchar` without having to consider the default character set of the server.

To use the `unichar` and `univarchar` datatypes with jConnect, you must perform the following two tasks:

- 1 Set the jConnect version to 6. See “Setting the jConnect version” on page 6 for more information.
- 2 Set the `DISABLE_UNICHAR_SENDING` connection property to false. See “Setting connection properties” on page 11 for more information.

When the server accepts `unichar` and `univarchar` data, jConnect will behave as follows:

- For all character data that a client wishes to send to the server – for example using the `PreparedStatement.setString(int column, String value)`– jConnect will determine if the string can be converted to the default character set of the server.
- If jConnect determines that the characters cannot be converted to the character set of the server (For example, some characters cannot be represented in the character set of the server), it will send the data to the server encoded as `unichar`/`univarchar` data.

For example, if a client attempted to send a unicode Japanese character to an Adaptive Server 12.5.1 which had `iso_1` as the default character set, `jConnect` would detect that the Japanese character could not be converted to an `iso_1` character. `jConnect` would then send the String as unicode data.

There is a performance penalty when a client sends `unichar/univarchar` data to a server. This is because `jConnect` must perform character-to-byte conversion twice for all Strings and characters that do not map directly to the default character set of the server. This penalty is why the `DISABLE_UNICHAR_SENDING` connection property defaults to `true`.

Note For more information on support for `unichar` and `univarchar` datatypes, see the manuals for Adaptive Server version 12.5 or later.

jConnect character-set converters

`jConnect` uses special classes for all character-set conversions. By selecting a character-set converter class, you specify how `jConnect` should handle single-byte and multibyte character-set conversions, and the performance impact the conversions will have on your applications.

There are two character-set conversion classes. The conversion class that `jConnect` uses is based on the version setting (for example, `VERSION_4`), and the `CHARSET` and `CHARSET_CONVERTER_CLASS` connection properties.

- The `TruncationConverter` class works only with single-byte character sets that use ASCII characters such as `iso_1` and `cp850`. It does not work with multibyte character sets or single-byte character sets that use non-ASCII characters.

Using the `TruncationConverter` class, `jConnect 5.x` handles character sets in the same manner as `jConnect` version 2.2. The `TruncationConverter` class is the default converter when the version setting is `VERSION_2`.

- The `PureConverter` class is a pure Java, multibyte character-set converter. `jConnect` uses this class if the version setting is `VERSION_4` or later. `jConnect` also uses this converter with `VERSION_2` if it detects a character set specified in the `CHARSET` connection property that is not compatible with the `TruncationConverter` class.

Although it enables multibyte character-set conversions, the `PureConverter` class may negatively impact `jConnect` driver performance. If driver performance is a concern, see “Improving character-set conversion performance” on page 36.

Selecting a character-set converter

`jConnect` uses the version setting from `SybDriver.setVersion()` to determine the default character-set converter class to use. For `VERSION_2`, the default is `TruncationConverter`. For `VERSION_4` and later, the default is `PureConverter`.

You can also set the `CHARSET_CONVERTER_CLASS` connection property to specify which character-set converter you want `jConnect` to use. This is useful if you want to use a character-set converter other than the default for your `jConnect` version.

For example, if you set `jConnect` to `VERSION_4` or later, but want to use the `TruncationConverter` class rather than the multibyte `PureConverter` class, you can set `CHARSET_CONVERTER_CLASS`:

```
...
props.put("CHARSET_CONVERTER_CLASS",
    "com.sybase.jdbc2.utils.TruncationConverter")
```

Setting the CHARSET connection property

You can specify the character set to use in your application by setting the `CHARSET` driver property. If you do not set the `CHARSET` property:

- For `VERSION_2`, `jConnect` uses `iso_1` as the default character set.
- For `VERSION_3`, `VERSION_4`, `VERSION_5`, and `VERSION_6`, `jConnect` uses the database’s default character set, and adjusts automatically to perform any necessary conversions on the client side.

You can also use the `-Jcharset` command line option for the `IsqlApp` application to specify a character set.

To determine which character sets are installed on your Adaptive Server, issue the following SQL query on your server:

```
select name from syscharsets
go
```

For the `PureConverter` class, if the designated `CHARSET` does not work with the client's Java Virtual Machine (VM), the connection fails with a `SQLException`, indicating that you must set `CHARSET` to a character set that is supported by both Adaptive Server and the client.

When the `TruncationConverter` class is used, character truncation is applied regardless of whether the designated `CHARSET` is 7-bit ASCII or not.

Improving character-set conversion performance

If you use multibyte character sets and need to improve driver performance, you can use the `SunloConverter` class provided with the `jConnect` samples. See “Character-set conversion” on page 113 for details.

Supported character sets

Table 2-4 lists the Sybase character sets that are supported by `jConnect`. The table also lists the corresponding JDK byte converter for each supported character set.

Although `jConnect` supports UCS-2, currently no Sybase databases or open servers support UCS-2.

Adaptive Server versions 12.5 and later support a version of Unicode known as the UTF-16 encoding.

Table 2-4 lists the character sets currently supported by Sybase.

Table 2-4: Supported Sybase character sets

SybCharset name	JDK byte converter
ascii_7	ASCII
big5	Big5
big5hk (see note)	Big5_HKSCS
cp037	Cp037
cp437	Cp437
cp500	Cp500
cp850	Cp850
cp852	Cp852
cp855	Cp855
cp857	Cp857
cp860	Cp860
cp863	Cp863
cp864	Cp864
cp866	Cp866
cp869	Cp869
cp874	Cp874
cp932	MS932
cp936	GBK
cp950	Cp950
cp1250	Cp1250
cp1251	Cp1251
cp1252	Cp1252
cp1253	Cp1253
cp1254	Cp1254
cp1255	Cp1255
cp1256	Cp1256
cp1257	Cp1257
cp1258	Cp1258
deckanji	EUC_JP
eucgb	EUC_CN
eucjis	EUC_JP
eucksc	EUC_KR
ibm420	Cp420
ibm918	Cp918
iso_1	ISO8859_1
iso88592	ISO8859-2

SybCharset name	JDK byte converter
is088595	ISO8859_5
iso88596	ISO8859_6
iso88597	ISO8859_7
iso88598	ISO8859_8
iso88599	ISO8859_9
iso15	ISO8859_15_FDIS
koi8	KOI8_R
mac	Macroman
mac_cyr	MacCyrillic
mac_ee	MacCentralEurope
macgreek	MacGreek
macturk	MacTurkish
sjis	MS932
tis620	MS874
utf8	UTF8

Note The big5hk character set will be supported only if you are using JDK 1.3 or later.

European currency symbol support

jConnect version 4.1 and later support the use of the new European currency symbol, or “*euro*” and its conversion to and from UCS-2 Unicode.

The *euro* has been added to the following Sybase character sets: cp1250, cp1251, cp1252, cp1253, cp1254, cp1255, cp1256, cp1257, cp1258, cp874, iso885915, and utf8.

To use the *euro* symbol:

- Use the PureConverter class, a pure Java, multibyte character-set converter. See “jConnect character-set converters” on page 34 for more information.
- Verify that the new character sets are installed on the server.

The *euro* symbol is currently supported only on Adaptive Server version 11.9.2 and later; Adaptive Server Anywhere version 8.0 and later provides support for the *euro* symbol.

- Select the appropriate character set on the client. See “Setting the CHARSET connection property” on page 35 for more information.

Unsupported character sets

The following Sybase character sets are not supported in jConnect 5.x because no JDK byte converters are analogous to the Sybase character sets:

- cp1047
- euccns
- greek8
- roman8
- turkish8

You can use these character sets with the `TruncationConverter` class as long as the application uses only the 7-bit ASCII subsets of these characters.

Working with databases

This section discusses database issues relevant to jConnect and includes these topics:

- Implementing high availability failover support
- Performing server-to-server remote procedure calls
- Wide table support for Adaptive Server version 12.5 and later
- Accessing database metadata
- Using cursors with result sets
- Support for batch updates
- Updating a database from a result set of a stored procedure
- Working with datatypes

Implementing high availability failover support

jConnect version 5.2 and later support the failover feature available in Adaptive Server version 12.0 and later.

Note Sybase failover in a high availability system is a different feature than “connection failover.” Sybase strongly recommends that you read this section *very carefully* if you want to use both.

Overview

Sybase failover allows you to configure two version 12.0 or later Adaptive Servers as companions. If the primary companion fails, that server’s devices, databases, and connections can be taken over by the secondary companion.

You can configure a high availability system either asymmetrically or symmetrically.

An *asymmetric* configuration includes two Adaptive Servers, each physically located on a different machine, that are connected so that if one of the servers is brought down, the other assumes its workload. The secondary Adaptive Server acts as a “hot standby” and does not perform any work until failover occurs.

A *symmetric* configuration also includes two Adaptive Servers running on separate machines. However, if failover occurs, either Adaptive Server can act as a primary or secondary companion for the other Adaptive Server. In this configuration, each Adaptive Server is fully functional with its own system devices, system databases, user databases, and user logins.

In either setup, the two machines are configured for dual access, which makes the disks visible and accessible to both machines.

You can enable failover in jConnect and connect a client application to an Adaptive Server configured for failover. If the primary server fails over to the secondary server, the client application also automatically switches to the second server and reestablishes network connections.

Note Refer to *Using Sybase Failover in High Availability Systems Manual for Adaptive Server* for more detailed information.

Requirements, dependencies, and restrictions

- You must have two version 12.0 or later Adaptive Servers configured for failover.
- You must use jConnect version 5.2, or later. Earlier driver versions do not support this feature.
- Only changes that were committed to the database before failover are retained when the client fails over.
- You must set the REQUEST_HA_SESSION jConnect connection property to true (see “Setting connection properties” on page 11).
- jConnect event notification does not work when failover occurs. See “Using event notification” on page 69.
- Close all statements when they are no longer used. jConnect stores information on statements to enable failover. If you do not close statements, you will experience memory leaks.

Implementing failover in jConnect

To implement failover support in jConnect choose one of the following two methods:

- 1 Use the two connection properties, REQUEST_HA_SESSION and SECONDARY_SERVER_HOSTPORT and set the following:
 - set REQUEST_HA_SESSION to *true*
 - set the SECONDARY_SERVER_HOSTPORT to the host name and port number where your secondary server is listening (see “Setting connection properties” on page 11, and the 'SECONDARY_SERVER_HOSTPORT' connection property in Table 2-2).
- 2 Use JNDI to connect to the server (See “Connecting to a server using JNDI”). Include an entry for the primary server and a separate entry for the secondary server in the directory service information file required by JNDI. The primary server entry will have an attribute (the HA OID) that refers to the entry for the secondary server.

Using LDAP as the service provider for JNDI, there are three possible forms that this HA attribute can have:

- *Relative Distinguished Name (RDN)* – this form assumes that the search base (typically provided by the `java.naming.provider.url` attribute) combined with this attribute's value is enough to identify the secondary server. For example, assume the primary server is at `hostname:4200` and the secondary server is at `hostname:4202`:

```
dn: servername=happrimary, o=Sybase, c=US
1.3.6.1.4.1.897.4.2.5: TCP#1#hostname 4200
1.3.6.1.4.1.897.4.2.15: servername=hasecondary
objectclass: sybaseServer
```

```
dn: servername=hasecondary, o=Sybase, c=US
1.3.6.1.4.1.897.4.2.5: TCP#1#hostname 4202
objectclass: sybaseServer
```

- *Distinguished Name (DN)* – this form assumes that the HA attribute's value uniquely identifies the secondary server, and may or may not duplicate values found in the search base. For example:

```
dn: servername=happrimary, o=Sybase, c=US
1.3.6.1.4.1.897.4.2.5: TCP#1#hostname 4200
1.3.6.1.4.1.897.4.2.15: servername=hasecondary,
                        o=Sybase, c=US ou=Accounting
objectclass: sybaseServer
```

```
dn: servername=hasecondary, o=Sybase, c=US, ou=Accounting
1.3.6.1.4.1.897.4.2.5: TCP#1#hostname 4202
objectclass: sybaseServer
```

Notice that `hasecondary` is located in a different branch of the tree (see the additional `ou=Accounting` qualifier).

- *Full LDAP URL* – this form assumes nothing about the search base. The HA attribute is expected to be a fully-qualified LDAP URL that is used to identify the secondary (it may even point to a different LDAP server). For example:

```

dn: servername=hafailover, o=Sybase, c=US
1.3.6.1.4.1.897.4.2.5: TCP#1#hostname 4200
1.3.6.1.4.1.897.4.2.15:
ldap://ldapserver: 386/servername=secondary,
    o=Sybase, c=US ou=Accounting
objectclass: sybaseServer

dn: servername=secondary, o=Sybase, c=US, ou=Accounting
1.3.6.1.4.1.897.4.2.5: TCP#1#hostname 4202
objectclass: sybaseServer

```

- In the directory service information file required by JNDI, set the REQUEST_HA_SESSION connection property to true to enable a failover session every time you make a connection.

The REQUEST_HA_SESSION connection property is used to indicate that the connecting client wants to begin a failover session with Adaptive Server version 12.0 or later, that is configured for failover. Setting this property to true causes jConnect to attempt a failover login. If you do not set this connection property, a failover session will not start, even if the server is configured correctly. The default value for REQUEST_HA_SESSION is false.

Set the connection property like any other connection property. You cannot reset the property once a connection has been made.

If you want more flexibility for requesting failover sessions, code the client application to set REQUEST_HA_SESSION at runtime.

The following example shows connection information entered for the database server SYBASE11 under an LDAP directory service:

```

dn: servername=SYBASE11, o=MyCompany, c=US
1.3.6.1.4.1.897.4.2.5: TCP#1#tahiti 3456
1.3.6.1.4.1.897.4.2.10: REPEAT_READ=false&PACKETSIZE=1024
1.3.6.1.4.1.897.4.2.10: CONNECTION_FAILOVER=false
1.3.6.1.4.1.897.4.2.11: pubs2
1.3.6.1.4.1.897.4.2.9: Tds
1.3.6.1.4.1.897.4.2.15: servername=SECONDARY
1.3.6.1.4.1.897.4.2.10: REQUEST_HA_SESSION=true

dn: servername=SECONDARY, o=MyCompany, c=US
1.3.6.1.4.1.897.4.2.5: TCP#1#moorea 6000

```

where “tahiti” is the primary server and “moorea” is the secondary companion server.

3 Request a connection using JNDI and LDAP.

- jConnect uses the LDAP server's directory server to determine the name and location of the primary and secondary servers:

```
/* get the connection */
Connection con = DriverManager.getConnection
("jdbc:sybase:jndi:ldap://ldap_server1:983" +
"/servername=Sybase11,o=MyCompany,c=US", props);

or
```

- Specify a searchbase:

```
props.put(Context.PROVIDER_URL,
"ldap://ldap_server1:983/ o=MyCompany, c=US");

Connection con=DriverManager.getConnection
("jdbc:sybase:jndi:servername=Sybase11", props);
```

Logging in to the primary server

If an Adaptive Server is not configured for failover, or for some reason cannot grant a failover session, the client cannot log in, and the following warning displays:

```
'The server denied your request to use the high-
availability feature.
```

```
Please reconfigure your database, or do not request a
high-availability session.'
```

Failing over to the secondary server

When failover occurs, the SQL exception JZ0F2 is thrown:

```
'Sybase high-availability failover has occurred. The
current transaction is aborted, but the connection is
still usable. Retry your transaction.'
```

The client then automatically reconnects to the secondary database using JNDI.

Note that:

- The identity of the database to which the client was connected and any committed transactions are retained.
- Partially read result sets, cursors, and stored procedure invocations are lost.
- When failover occurs, your application may need to restart a procedure or go back to the last completed transaction or activity.

Failing back to the primary server

At some point, the client will fail back from the secondary server to the primary server. When failback occurs is determined by the System Administrator who issues `sp_failback` on the secondary server. Afterward, the client can expect the same behavior and results on the primary server as documented in “Failing over to the secondary server” on page 44.

Performing server-to-server remote procedure calls

A Transact-SQL language command or stored procedure running on one server can execute a stored procedure located on another server. The server to which an application has connected logs in to the remote server, and executes a server-to-server remote procedure call.

An application can specify a “universal” password for server-to-server communication; that is, a password used in all server-to-server connections. Once the connection is open, the server uses this password to log in to any remote server. By default, `jConnect` uses the current connection’s password as the default password for server-to-server communications.

However, if the passwords are different on two servers for the same user and that user is performing server-to-server remote procedure calls, the application must explicitly define passwords for each server it plans to use.

`jConnect` versions 4.1 and later include a property that lets you set a universal “remote” password or different passwords on several servers. `jConnect` lets you set and configure the property using the `setRemotePassword()` method in the `SybDriver` class:

```
Properties connectionProps = new Properties();

public final void setRemotePassword(String serverName,
    String password, Properties connectionProps
```

To use this method, the application needs to import the `SybDriver` class, then call the method:

```
import com.sybase.jdbcx.SybDriver;
SybDriver sybDriver = (SybDriver)
    Class.forName("com.sybase.jdbc2.jdbc.SybDriver").n
        ewInstance();
sybDriver.setRemotePassword
    (serverName, password, connectionProps);
```

Note To set different remote passwords for various servers, repeat the preceding call (appropriate for your version of jConnect) for each server.

This call adds the given server name-password pair to the given Properties object, which can be passed by the application to DriverManager in `DriverManager.getConnection(server_url, props)`.

If `serverName` is NULL, the universal password will be set to `password` for subsequent connections to all servers except the ones specifically defined by previous calls to `setRemotePassword()`.

When an application sets the REMOTEPWD property, jConnect no longer sets the default universal password.

Wide table support for Adaptive Server version 12.5 and later

Adaptive Server version 12.5 and later offers limits and parameters that are larger than in previous versions of the database server. For example:

Note You must use jConnect 5.5 or later to access the wide table functionality in Adaptive Server version 12.5 and later.

- Tables can contain 1,024 columns.
- Varchar and varbinary columns can contain more than 255 bytes of data.
- You can send and retrieve up to 2,048 parameters when invoking stored procedures or inserting data into tables.

To take advantage of this capability, jConnect version 5.5 and later users need to set their JCONNECT_VERSION property to VERSION_6 or VERSION_LATEST. This will request that the server enable wide table support.

Note jConnect continues to work with an Adaptive Server version 12.5 and later if you set the version to below VERSION_6. However, if you try selecting from a table that requires wide table support to fully retrieve the data, you may encounter unexpected errors or data truncation.

You can also set the version to VERSION_6 or VERSION_LATEST when you access data from a Sybase server that does not support wide tables. In this case, the server simply ignores your request for wide table support.

Wide table support offers an extra benefit for jConnect users, besides the larger number of columns and parameters—a greater amount of `ResultSetMetaData`. For example, in versions of jConnect earlier than 5.5, the `ResultSetMetaData` methods `getCatalogName`, `getSchemaName`, and `getTableName` all returned “Not Implemented” `SQLExceptions` because that metadata was not supplied by the server. When you enable wide table support, the server now sends back this information, and the three methods return useful information.

Accessing database metadata

To support JDBC `DatabaseMetaData` methods, Sybase provides a set of stored procedures that jConnect can call for metadata about a database. These stored procedures must be installed on the server for the JDBC metadata methods to work.

If the stored procedures for providing metadata are not already installed in a Sybase server, you can install them using stored procedure scripts provided with jConnect:

- *sql_server.sql* installs stored procedures on pre-12.0 Adaptive Server databases.
- *sql_server12.sql* installs stored procedures on a version 12.0 Adaptive Server database.
- *sql_server12.5.sql* installs stored procedures on Adaptive Server databases of version 12.5 and later.

- *sql_asa.sql* installs stored procedures on an Adaptive Server Anywhere database.

Note The most recent version of these scripts is compatible with all versions of jConnect.

See the *Sybase jConnect for JDBC Installation Guide* and *Release Bulletin* for complete instructions on installing stored procedures.

In addition, to use the metadata methods, you must set the `USE_METADATA` connection property to true (its default value) when you establish a connection.

You cannot get metadata about temporary tables in a database.

Note The `DatabaseMetaData.getPrimaryKeys()` method finds primary keys declared in a table definition (`CREATE TABLE`) or with alter table (`ALTER TABLE ADD CONSTRAINT`). It does not find keys defined using `sp_primarykey`.

Server-side metadata installation

Metadata support can be implemented in either the client (ODBC, JDBC) or in the data source (server stored procedures). jConnect provides metadata support on the server, which results in the following benefits:

- Maintains jConnect's small size, which ensures the driver can be quickly downloaded from the Internet.
- Gains runtime efficiency from preloaded stored procedures on the data source.
- Provides flexibility—jConnect can connect to a variety of databases.

Using cursors with result sets

jConnect 5.x implements many JDBC 2.0 cursor and update methods. These methods make it easier to use cursors and to update rows in a table based on values in a result set.

In JDBC 2.0, `ResultSets` are characterized by their type and their concurrency. The type and concurrency values are part of the `java.sql.ResultSet` interface and are described in its javadocs.

Table 2-5 identifies the characteristics of `java.sql.ResultSet` that are available in `jConnect 5.x`.

Table 2-5: *java.sql.ResultSet* options available in *jConnect 5.x*

Concurrency	Type		
	TYPE_FORWARD_ONLY	TYPE_SCROLL_INSENSITIVE	TYPE_SCROLL_SENSITIVE
<i>CONCUR_READ_ONLY</i>	Supported in 5.x	Supported in 5.x	Not available in 5.x
<i>CONCUR_UPDATABLE</i>	Supported in 5.x	Not available in 5.x	Not available in 5.x

This section includes the following topics:

- Creating a cursor
- Positioned updates and deletes using JDBC 1.x methods
- Using a cursor with a `PreparedStatement` object
- Support for `SCROLL_INSENSITIVE` result sets in `jConnect`

Creating a cursor

To create a cursor using `jConnect` there are two methods:

- `SybStatement.setCursorName()`

Use `SybStatement.setCursorName()`, to explicitly assign the cursor a name. The signature for `SybStatement.setCursorName()` is:

```
void setCursorName(String name) throws SQLException;
```

- `SybStatement.setFetchSize()`

Use `SybStatement.setFetchSize()` to create a cursor and specify the number of rows returned from the database in each fetch. The signature for `SybStatement.setFetchSize()` is:

```
void setFetchSize(int rows) throws SQLException;
```

When you use `setFetchSize()` to create a cursor, the `jConnect` driver names the cursor. To get the cursor's name, use `ResultSet.setCursorName()`.

Another way you can create cursors is to specify the kind of `ResultSet` you want returned by the statement, using the following JDBC 2.0 method on the connection:

```
Statement createStatement(int resultSetType, int
resultSetConcurrency) throws SQL Exception
```

The type and concurrencies correspond to the types and concurrencies found on the `ResultSet` interface listed in Table 2-5. If you request an unsupported `ResultSet`, a SQL warning is chained to the connection. When the returned `Statement` is executed, you will receive the kind of `ResultSet` that is most like the one you requested. See the JDBC 2.0 specification for more details on this method's behavior.

If you do not use `createStatement()`, the default types of `ResultSet` are:

- If you call only `Statement.executeQuery()`, then the `ResultSet` returned is a `SybResultSet` that is `TYPE_FORWARD_ONLY` and `CONCUR_READ_ONLY`.
- If you call `setFetchSize()` or `setCursorName()`, then the `ResultSet` returned from `executeQuery()` is a `SybCursorResultSet` that is `TYPE_FORWARD_ONLY` and `CONCUR_UPDATABLE`.

To verify that the kind of `ResultSet` object is what you intended, use the following two `ResultSet` methods:

```
int getConcurrency() throws SQLException;  
int getType() throws SQLException;
```

The basic steps for creating and using a cursor are:

- 1 Create the cursor using `Statement.setCursorName()` or `SybStatement.setFetchSize()`.
- 2 Invoke `Statement.executeQuery()` to open the cursor for a statement and return a cursor result set.
- 3 Invoke `ResultSet.next()` to fetch rows and position the cursor in the result set.

The following example uses each of the two methods for creating cursors and returning a result set. It also uses `ResultSet.setCursorName()` to get the name of the cursor created by `SybStatement.setFetchSize()`.

```
// With conn as a Connection object, create a  
// Statement object and assign it a cursor using  
// Statement.setCursorName().  
Statement stmt = conn.createStatement();  
stmt.setCursorName("author_cursor");  
  
// Use the statement to execute a query and return  
// a cursor result set.  
ResultSet rs = stmt.executeQuery("SELECT au_id,  
    au_lname, au_fname FROM authors  
    WHERE city = 'Oakland'");
```

```
while(rs.next())
{
    ...
}

// Create a second statement object and use
// SybStatement.setFetchSize() to create a cursor
// that returns 10 rows at a time.
SybStatement syb_stmt = conn.createStatement();
syb_stmt.setFetchSize(10);

// Use the syb_stmt to execute a query and return
// a cursor result set.
SybCursorResultSet rs2 =
    (SybCursorResultSet)syb_stmt.executeQuery
    ("SELECT au_id, au_lname, au_fname FROM
authors
    WHERE city = 'Pinoles'");
while(rs2.next())
{
    ...
}

// Get the name of the cursor created through the
// setFetchSize() method.
String cursor_name = rs2.getCursorName();
...
// For JDBC 5.x, create a third statement
// object using the new method on Connection,
// and obtain a SCROLL_INSENSITIVE ResultSet.
// Note: you no longer have to downcast the
// Statement or the ResultSet.
Statement stmt = conn.createStatement(
    ResultSet.TYPE_SCROLL_INSENSITIVE,
    ResultSet.CONCUR_READ_ONLY);
ResultSet rs3 = stmt.executeQuery
    ("SELECT ... [whatever]");
// Execute any of the JDBC 2.0 methods that
// are valid for read only ResultSets.
rs3.next();
rs3.previous();
rs3.relative(3);
rs3.afterLast();
...
```

Positioned updates and deletes using JDBC 1.x methods

The following example shows how to use methods in JDBC 1.x to do a positioned update. The example creates two `Statement` objects, one for selecting rows into a cursor result set, and the other for updating the database from rows in the result set.

Note Although this manual provides sample code relating to JDBC 1.0 and 2.0 methods, Sybase strongly suggests that you use JDBC 2.0 for ease of use and portability.

```
// Create two statement objects and create a cursor
// for the result set returned by the first
// statement, stmt1. Use stmt1 to execute a query
// and return a cursor result set.
Statement stmt1 = conn.createStatement();
Statement stmt2 = conn.createStatement();
stmt1.setCursorName("author_cursor");
ResultSet rs = stmt1.executeQuery("SELECT
    au_id, au_lname, au_fname
    FROM authors WHERE city = 'Oakland'
    FOR UPDATE OF au_lname");

// Get the name of the cursor created for stmt1 so
// that it can be used with stmt2.
String cursor = rs.getCursorName();

// Use stmt2 to update the database from the
// result set returned by stmt1.
String last_name = new String("Smith");
while(rs.next())
{
    if (rs.getString(1).equals("274-80-9391"))
    {
        stmt2.executeUpdate("UPDATE authors "+
            "SET au_lname = "+last_name +
            "WHERE CURRENT OF " + cursor);
    }
}
```

Deletions in a result set

The following example uses `Statement` object `stmt2`, from the preceding code, to perform a positioned deletion:

```
stmt2.executeUpdate("DELETE FROM authors  
WHERE CURRENT OF " + cursor);
```

Positioned updates and deletes using JDBC 2.0 methods

This section discusses JDBC 2.0 methods for updating columns in the current cursor row and updating the database from the current cursor row in a result set. They are followed by an example.

Updating columns in a result set

JDBC 2.0 specifies a number of methods for updating column values from a result set in memory, on the client. The updated values can then be used to perform an update, insert, or delete operation on the underlying database. All of these methods are implemented in the `SybCursorResultSet` class.

Examples of some of the JDBC 2.0 update methods available in `jConnect` are:

```
void updateAsciiStream(String columnName, java.io.InputStream x,  
    int length) throws SQLException;  
void updateBoolean(int columnIndex, boolean x) throws  
    SQLException;  
void updateFloat(int columnIndex, float x) throws SQLException;  
void updateInt(String columnName, int x) throws SQLException;  
void updateInt(int columnIndex, int x) throws SQLException;  
void updateObject(String columnName, Object x) throws  
    SQLException;
```

Methods for updating the database from a result set

JDBC 2.0 specifies two new methods for updating or deleting rows in the database, based on the current values in a result set. These methods are simpler in form than `Statement.executeUpdate()` in JDBC 1.x and do not require a cursor name. They are implemented in `SybCursorResultSet`:

```
void updateRow() throws SQLException;  
void deleteRow() throws SQLException;
```

Note The concurrency of the result set must be `CONCUR_UPDATABLE`, otherwise the above methods will raise an exception. For `insertRow()`, all table columns that require non-null entries must be specified.

Methods provided on `DatabaseMetaData` dictate when these changes are visible.

Example

The following example creates a single Statement object that is used to return a cursor result set. For each row in the result set, column values are updated in memory and then the database is updated with the row's new column values.

```
// Create a Statement object and set fetch size to
// 25. This creates a cursor for the Statement
// object Use the statement to return a cursor
// result set.
SybStatement syb_stmt =
(SybStatement)conn.createStatement();
syb_stmt.setFetchSize(25);
SybCursorResultSet syb_rs =
(SybCursorResultSet)syb_stmt.executeQuery(
    "SELECT * from T1 WHERE ...")

// Update each row in the result set according to
// code in the following while loop. jConnect
// fetches 25 rows at a time, until fewer than 25
// rows are left. Its last fetch takes any
// remaining rows.
while(syb_rs.next())
{
    // Update columns 2 and 3 of each row, where
    // column 2 is a varchar in the database and
    // column 3 is an integer.
    syb_rs.updateString(2, "xyz");
    syb_rs.updateInt(3,100);
    //Now, update the row in the database.
    syb_rs.updateRow();
}
// Create a Statement object using the
// JDBC 2.0 method implemented in jConnect 5.x
Statement stmt = conn.createStatement
(ResultSet.TYPE_FORWARD_ONLY, ResultSet.CONCUR_UPDATABLE);
// In jConnect 5.x, downcasting to SybCursorResultSet is not
// necessary. Update each row in the ResultSet in the same
// manner as above
while (rs.next())
{
    rs.updateString(2, "xyz");
    rs.updateInt(3,100);
    rs.updateRow();
// Use the Statement to return an updatable ResultSet
ResultSet rs = stmt.executeQuery("SELECT * FROM T1 WHERE...");
}
```


Deleting a row from a ResultSet

To delete a row from a cursor result set, you can use `SybCursorResultSet.deleteRow()` as follows:

```
while(syb_rs.next())
{
    int col3 = getInt(3);
    if (col3 >100)
    {
        syb_rs.deleteRow();
    }
}
```

Inserting a row into a ResultSet

The following example illustrates how to do inserts using the JDBC 2.0 API. There is no need to downcast to a `SybCursorResultSet`.

```
// prepare to insert
rs.moveToInsertRow();
// populate new row with column values
rs.updateString(1, "New entry for col 1");
rs.updateInt(2, 42);
// insert new row into db
rs.insertRow();
// return to current row in result set
rs.moveToCurrentRow();
```

Using a cursor with a PreparedStatement object

Once you create a `PreparedStatement` object, you can use it multiple times with the same or different values for its input parameters. If you use a cursor with a `PreparedStatement` object, you need to close the cursor after each use and then reopen the cursor to use it again. A cursor is closed when you close its result set (`ResultSet.close()`). It is opened when you execute its prepared statement (`PreparedStatement.executeQuery()`).

The following example shows how to create a `PreparedStatement` object, assign it a cursor, and execute the `PreparedStatement` object twice, closing and then reopening the cursor.

```
// Create a prepared statement object with a
// parameterized query.
PreparedStatement prep_stmt =
conn.prepareStatement(
"SELECT au_id, au_lname, au_fname "+
```

```
"FROM authors WHERE city = ? "+
"FOR UPDATE OF au_lname");

//Create a cursor for the statement.
prep_stmt.setCursorName("author_cursor");

// Assign the parameter in the query a value.
// Execute the prepared statement to return a
// result set.
prep_stmt.setString(1, "Oakland");
ResultSet rs = prep_stmt.executeQuery();

//Do some processing on the result set.
while(rs.next())
{
    ...
}

// Close the result, which also closes the cursor.
rs.close();

// Execute the prepared statement again with a new
// parameter value.
prep_stmt.setString(1,"San Francisco");
rs = prep_stmt.executeQuery();
// reopens cursor
```

Support for SCROLL_INSENSITIVE result sets in jConnect

jConnect version 5.x supports only TYPE_SCROLL_INSENSITIVE result sets.

jConnect uses the Tabular Data Stream (TDS)—Sybase’s proprietary protocol—to communicate with Sybase database servers. As of jConnect 5.x, TDS does not support scrollable cursors. To support scrollable cursors, jConnect caches the row data on demand, on the client, on each call to `ResultSet.next()`. However, when the end of the result set is reached, the entire result set is stored in the client’s memory. Because this may cause a performance strain, we recommend that you use `TYPE_SCROLL_INSENSITIVE` result sets only when the result set is reasonably small.

Note When you use `TYPE_SCROLL_INSENSITIVE` `ResultSet`s in jConnect 5.x, you can only call the `isLast()` method after the last row of the `ResultSet` has been read. Calling `isLast()` before the last row is reached will cause an `UnimplementedOperationException` to be thrown.

jConnect provides the `ExtendResultSet` in the *sample2* directory; this sample provides a limited `TYPE_SCROLL_INSENSITIVE` `ResultSet` using JDBC 1.0 interfaces.

This implementation uses standard JDBC 1.0 methods to produce a scroll-insensitive, read-only result set; that is, a static view of the underlying data that is not sensitive to changes made while the result set is open. `ExtendedResultSet` caches all of the `ResultSet` rows on the client. Be cautious when you use this class with large result sets.

The `sample.ScrollableResultSet` interface:

- Is an extension of JDBC 1.0 `java.sql.ResultSet`.
- Defines additional methods that have the same signatures as the JDBC 2.0 `java.sql.ResultSet`.
- Does *not* contain all of the JDBC 2.0 methods. The missing methods deal with modifying the `ResultSet`.

The methods from the JDBC 2.0 API are:

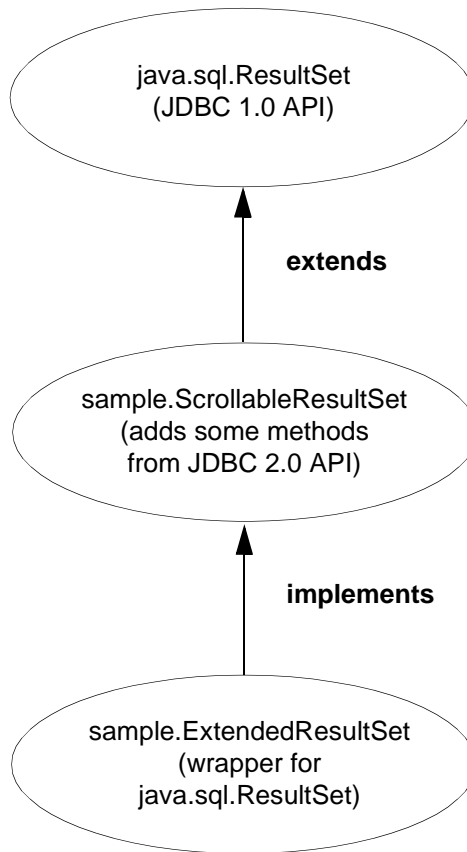
```
boolean previous() throws SQLException;
boolean absolute(int row) throws SQLException;
boolean relative(int rows) throws SQLException;
boolean first() throws SQLException;
boolean last() throws SQLException;
void beforeFirst() throws SQLException;
void afterLast() throws SQLException;
boolean isFirst() throws SQLException;
boolean isLast() throws SQLException;
```

```
boolean isBeforeFirst() throws SQLException;
boolean isAfterLast() throws SQLException;
int getFetchSize() throws SQLException;
void setFetchSize(int rows) throws SQLException;
int getFetchDirection() throws SQLException;
void setFetchDirection(int direction) throws
SQLException;
int getType() throws SQLException;
int getConcurrency() throws SQLException;
int getRow() throws SQLException;
```

To use the new sample classes, create an `ExtendedResultSet` using any JDBC 1.0 `java.sql.ResultSet`. Below are the relevant pieces of code (assume a Java 1.1 environment):

```
// import the sample files
import sample.*;
//import the JDBC 1.0 classes
import java.sql.*;
// connect to some db using some driver;
// create a statement and a query;
// Get a reference to a JDBC 1.0 ResultSet
ResultSet rs = stmt.executeQuery(_query);
// Create a ScrollableResultSet with it
ScrollableResultSet srs = new ExtendedResultSet(rs);
// invoke methods from the JDBC 2.0 API
srs.beforeFirst();
// or invoke methods from the JDBC 1.0 API
if (srs.next())
    String column1 = srs.getString(1);
```

Figure 2-1 is a class diagram that shows the relationships between the new sample classes and the JDBC API.

Figure 2-1: Class diagram

See the JDBC 2.0 API at <http://java.sun.com/products/jdbc/jdbcse2.html> for more details.

Support for batch updates

Batch updates allow a `Statement` object to submit multiple update commands as one unit (batch) to an underlying database for processing together.

Note To use batch updates, you must install the latest metadata scripts provided in the `sp` directory under your jConnect installation directory.

See *BatchUpdates.java* in the *sample2* subdirectories for an example of using batch updates with Statement, PreparedStatement, and CallableStatement.

jdbcConnect also supports dynamic PreparedStatements in batch.

Implementation notes

jdbcConnect implements batch updates as specified in the JDBC 2.0 API, except as described below.

- If the JDBC 2.0 standard for implementing `BatchUpdateException.getUpdateCounts()` is modified or relaxed in the future, jdbcConnect will continue to implement the original standard by having `BatchUpdateException.getUpdateCounts()` return an `int[]` length of $M < N$, indicating that the first M statements in the batch succeeded, that the $M+1$ statement failed, and $M+2..N$ statements were not executed; where “ N ” equals the total statements in the batch.
- Batch updates of stored procedures – to call stored procedures in batch (unchained) mode, you must create the stored procedure in unchained mode. For more information, see “Stored procedure executed in unchained transaction mode” on page 109.
- Adaptive Server version 11.5.x and later – if the server encounters an error during batch execution, `BatchUpdateException.getUpdateCounts()` will return only an `int[]` length of zero. The entire transaction is rolled back if an error is encountered, resulting in zero successful rows.

Note The transaction is **not** rolled back if the error is a duplicate key row insert (see next note for more information).

- Adaptive Server – a duplicate key row insertion does not result in the termination and rollback of batch statements. The server will continue processing the statements in the batch until you issue a cancel, or the batch completes or encounters an error, other than a duplicate key row insertion. Because jdbcConnect sends a cancel to the server when it detects any exception (including duplicate key row insertion) during batch processing, it is impossible to determine exactly how much of the batch the server executed before receiving the cancel. Therefore, Sybase strongly recommends that in accordance with the JDBC specification, you should execute batches inside of transactions with `autoCommit` set to `false`. In doing it this way, you can roll back your transactions and return the database to a known state before retrying the batch.

- Adaptive Server version 11.0.1 – returns 0 (zero) rows affected for stored procedures.
- SQL Anywhere version 5.5.x:
 - SQL Anywhere version 5.5.x does not allow you to obtain inserted row counts from stored procedures that contain inserts. For example:

```
create proc sp_A as insert tableA values (1,
'hello A')
create proc sp_B
as
    insert tableA values (1, 'hello A')
update tableA set coll=2
create proc sp_C
as
update tableA set coll=2
delete tableA
```

Running `executeBatch` on the preceding stored procedures would result in, respectively:

```
0 Rows Affected
1 Rows Affected
2 Rows Affected
```

- There is no support for dynamic `PreparedStatement`s in batch.
- Because SQL Anywhere 5.5.x does not natively support batch updates according to the JDBC 2.0 specification, batch updates are carried out in an `executeUpdate` loop.
- Batch updates in databases that do not support batch updates – `jConnect` carries out batch updates in an `executeUpdate` loop even if your database does not support batch updates. This allows you to use the same batch code, regardless of the database to which you are pointing.

See *Sun Microsystems, Inc. JDBC™ 2.0 API* for more details on batch updates.

Updating a database from a result set of a stored procedure

`jConnect` includes `update` and `delete` methods that allow you to get a cursor on the result set returned by a stored procedure. You can then use the cursor's position to update or delete rows in the underlying table that provided the result set. The methods are in `SybCursorResultSet`:

```
void updateRow(String tableName) throws SQLException;
```

`void deleteRow(String tableName) throws SQLException;`

The *tableName* parameter identifies the database table that provided the result set.

To get a cursor on the result set returned by a stored procedure, you need to use either `SybCallableStatement.setCursorName()` or `SybCallableStatement.setFetchSize()` before you execute the callable statement that contains the procedure. The following example shows how to create a cursor on the result set of a stored procedure, update values in the result set, and then update the underlying table using the `SybCursorResultSet.update()` method:

```
// Create a CallableStatement object for executing the stored
// procedure.
CallableStatement sproc_stmt =
    conn.prepareCall("{call update_titles}");

// Set the number of rows to be returned from the database with
// each fetch. This creates a cursor on the result set.
(SybCallableStatement)sproc_stmt.setFetchSize(10);

//Execute the stored procedure and get a result set from it.
SybCursorResultSet sproc_result = (SybCursorResultSet)
    sproc_stmt.executeQuery();

// Move through the result set row by row, updating values in the
// cursor's current row and updating the underlying titles table
// with the modified row values.
while(sproc_result.next())
{
    sproc_result.updateString(...);
    sproc_result.updateInt(...);
    ...
    sproc_result.updateRow(titles);
}
```

Working with datatypes

Sending numeric data

`jConnect` has added the `SybPreparedStatement` extension to support the way Adaptive Server handles the `NUMERIC` datatype where precision (total digits) and scale (digits after the decimal) can be specified.

The corresponding datatype in Java—`java.math.BigDecimal`—is slightly different, and these differences can cause problems when `jConnect` applications use the `setBigDecimal` method to control values of an input/output parameter. Specifically, there are cases where the precision and scale of the parameter must precisely match that precision and scale of the corresponding SQL object, whether it is a stored procedure parameter or a column.

To give `jConnect` applications fuller control over the `setBigDecimal` method, The `SybPreparedStatement` extension has been added with this method:

```
public void setBigDecimal (int parameterIndex, BigDecimal X, int scale,  
    int precision) throws SQLException
```

See the `SybPrepExtension.java` sample in the `/sample2` subdirectories under your `jConnect` installation directory for more information.

Updating image data in the database

`jConnect` has a `TextPointer` class with `sendData()` methods for updating an *image* column in an Adaptive Server or Adaptive Server Anywhere database. In earlier versions of `jConnect`, you had to send image data using the `setBinaryStream()` method in `java.sql.PreparedStatement`. The `TextPointer.sendData()` methods use `java.io.InputStream` and greatly improve performance when you send image data to an Adaptive Server database.

Warning! The `TextPointer` class has been deprecated; that is, it is no longer recommended and may cease to exist in a future version of `jConnect`.

If your data server is Adaptive Server version 12.5 or later or Adaptive Server Anywhere version 6.0 or later, use the standard JDBC form to send image data:

```
PreparedStatement.setBinaryStream(int paramIndex,  
    InputStream image)
```

To obtain instances of the `TextPointer` class, you can use either of two `getTextPtr()` methods in `SybResultSet`:

```
public TextPointer getTextPtr(String columnName)  
public TextPointer getTextPtr(int columnIndex)
```

Public methods in the `TextPointer` class

The `com.sybase.jdbc` package contains the `TextPointer` class. Its public method interface is:

```
public void sendData(InputStream is, boolean log)
    throws SQLException
```

```
public void sendData(InputStream is, int length,
    boolean log) throws SQLException
```

```
public void sendData(InputStream is, int offset,
    int length, boolean log) throws SQLException
```

```
public void sendData(byte[] byteInput, int offset,
    int length, boolean log) throws SQLException
```

`sendData(InputStream is, boolean log)` – Updates an image column with data in the specified input stream.

`sendData(InputStream is, int length, boolean log)` – updates an image column with data in the specified input stream. *length* is the number of bytes being sent.

`sendData(InputStream is, int offset, int length, boolean log)` – updates an image column with data in the specified input stream, starting at the byte offset given in the *offset* parameter and continuing for the number of bytes specified in the *length* parameter.

`sendData(byte[] byteInput, int offset, int length, boolean log)` – updates a column with image data contained in the byte array specified in the *byteInput* parameter. The update starts at the byte offset given in the *offset* parameter and continues for the number of bytes specified in the *length* parameter.

Each method has a *log* parameter. The *log* parameter specifies whether image data is to be fully logged in the database transaction log. If the *log* parameter is set to true, the entire binary image is written into the transaction log. If the *log* parameter is set to false, the update is logged, but the image itself is not included in the log.

❖ Updating an *image* column with *TextPointer.sendData()*

To update a column with image data:

- 1 Get a *TextPointer* object for the row and column that you want to update.
- 2 Use *TextPointer.sendData()* to execute the update.

The next two sections illustrate these steps with an example. In the example, image data from the file *Anne_Ringer.gif* is sent to update the *pic* column of the *au_pix* table in the *pubs2* database. The update is for the row with author ID 899-46-2035.

Getting a *TextPointer*
object

text and image columns contain timestamp and page-location information that is separate from their text and image data. When data is selected from a text or image column, this extra information is “hidden” as part of the result set.

A `TextPointer` object for updating an image column requires this hidden information, but does not need the image portion of the column data. To get this information, you need to select the column into a `ResultSet` object and then use `SybResultSet.getTextPtr()` (see the example that follows the next paragraph). `SybResultSet.getTextPtr()` extracts text-pointer information, ignores image data, and creates a `TextPointer` object.

When a column contains a significant amount of image data, selecting the column for one or more rows and waiting to get all the data is likely to be inefficient, since the data is not used. You can shortcut this process by using the `set textsize` command to minimize the amount of data returned in a packet. The following code example for getting a `TextPointer` object includes the use of `set textsize` for this purpose.

```
/*
 * Define a string for selecting pic column data for author ID
 * 899-46-2035.
 */
String getColumnData = "select pic from au_pix where au_id = '899-46-2035'";

/*
 * Use set textsize to return only a single byte of column data
 * to a Statement object. The packet with the column data will
 * contain the "hidden" information necessary for creating a
 * TextPointer object.
 */
Statement stmt= connection.createStatement();
stmt.executeUpdate("set textsize 1");

/*
 * Select the column data into a ResultSet object--cast the
 * ResultSet to SybResultSet because the getTextPtr method is
 * in SybResultSet, which extends ResultSet.
 */
SybResultSet rs = (SybResultSet)stmt.executeQuery(getColumnData);

/*
 * Position the result set cursor on the returned column data
 * and create the desired TextPointer object.
 */
rs.next();
TextPointer tp = rs.getTextPtr("pic");

/*
 * Now, assuming we are only updating one row, and won't need
```

```
* the minimum textsize set for the next return from the server,  
* we reset textsize to its default value.  
*/  
stmt.executeUpdate("set textsize 0");
```

Executing the
update with
TextPointer.sendData

The following code uses the *TextPointer* object from the preceding section to update the *pic* column with image data in the file *Anne_Ringer.gif*.

```
/*  
 *First, define an input stream for the file.  
 */  
FileInputStream in = new FileInputStream("Anne_Ringer.gif");  
  
/*  
 * Prepare to send the input stream without logging the image data  
 * in the transaction log.  
 */  
boolean log = false;  
  
/*  
 * Send the image data in Anne_Ringer.gif to update the pic  
 * column for author ID 899-46-2035.  
 */  
tp.sendData(in, log);
```

See the *TextPointers.java* sample in the *sample2* subdirectories under your jConnect installation directory for more information.

Using text data

In earlier versions, jConnect used a *TextPointer* class with *sendData()* methods for updating a *text* column in an Adaptive Server or Adaptive Server Anywhere database.

The *TextPointer* class has been deprecated; that is, it is no longer recommended and may cease to exist in a future version of Java.

If your data server is Adaptive Server 12.5 or later or Adaptive Server Anywhere version 6.0 or later, use the standard JDBC form to send text data:

```
PreparedStatement.setAsciiStream(int paramIndex,  
    InputStream text, int length)
```

or

```
PreparedStatement.setUnicodeStream(int paramIndex,  
    InputStream text, int length)
```

or

```
PreparedStatement.setCharacterStream(int paramIndex,  
    Reader reader, int length)
```

Using *Date* and *Time* datatypes

Adaptive Server versions 12.5.1 and later offer support for the SQL date and time datatypes. Previously, Adaptive Server offered only support for the `datetime` and `smalldatetime` datatypes. These datatypes were limited for the following reasons:

- There was not a way to have separate time and date datatypes.
- Dates prior to 1/1/1753 were illegal. Datetime values could hold only dates between 1/1/1753 and 12/31/9999.
- When JDBC clients used the `setTime` and `setDate` methods to insert `java.sql.Time` or `java.sql.Date` methods, the values were converted to `datetime` datatypes in the server. This conversion resulted in the addition of a default date or time to the insert value supplied by the client.

The addition of the date and time datatypes provides the following advantages:

- Date values can now be between Jan. 1, 0001 and Dec. 31, 9999, exactly matching the allowable values in `java.sql.Date`.
- A direct mapping now exists between `java.sql.Date` and the date datatype, as well as between `java.sql.Time` and the time datatype.

The date and time datatypes can be used only with jConnect 5.5 and later. To use them, a jConnect client must set the driver version to 6. To do so, the client can use the `SybDriver.setVersion()` method, or use the `JCONNECT_VERSION` connection property.

Implementation notes

- You will not have to change any of the application code if you decide to use the date and time datatypes by setting the version to 6.

- If you select from a table that contains a date or time column, and you have not enabled date/time support in jConnect (by setting the version), the server will try to convert the date or time to a datetime value before returning it. This can cause problems if the date to be returned is prior to 1/1/1753. In that case, a conversion error will occur, and the database will inform you of the error.
- Adaptive Server Anywhere supports a date and time datatype, but the date and time datatypes are not yet directly compatible with those in Adaptive Server version 12.5.1 and later. Using jConnect, you should continue to use the datetime and smalldatetime datatypes when communicating with Adaptive Server Anywhere.
- The maximum value in a datetime column in Adaptive Server Anywhere is 1-1-7911 00:00:00.
- Using jConnect you will receive conversion errors if you attempt to insert dates prior to 1/1/1753 into datetime columns or parameters
- Refer to the Adaptive Server manuals for more information on the date and time datatypes; of special note is the section on allowable implicit conversions.
- If you use getObject() with an Adaptive Server date, time, or datetime column, the value returned will be, respectively, a java.sql.Date, java.sql.Time, or java.sql.Timestamp datatype.

Char/Varchar/Text datatypes and getByte()

Do not use rs.getByte() on a char, univarchar, unichar, varchar, or text field unless the data is hex, octal, or decimal.

Implementing advanced features

This section describes how to use advanced jConnect features and contains the following topics:

- Using event notification
- Handling error messages
- Storing Java objects as column data in a table
- Dynamic class loading

- JDBC 2.0 optional package extensions support

Using event notification

You can use the jConnect event notification feature to have your application notified when an Open Server procedure is executed.

To use this feature, you must use the `SybConnection` class, which extends the `Connection` interface. `SybConnection` contains a `regWatch()` method for turning event notification on and a `regNoWatch()` method for turning event notification off.

Your application must also implement the `SybEventHandler` interface. This interface contains one public method, `void event(String proc_name, ResultSet params)`, which is called when the specified event occurs. The parameters of the event are passed to `event()` and it tells the application how to respond.

To use event notification in your application, call `SybConnection.regWatch()` to register your application in the notification list of a registered procedure. Use this syntax:

```
SybConnection.regWatch(proc_name,eventHdlr,option)
```

- *proc_name* is a `String` that is the name of the registered procedure that generates the notification.
- *eventHdlr* is an instance of the `SybEventHandler` class that you implement.
- *option* is either `NOTIFY_ONCE` or `NOTIFY_ALWAYS`. Use `NOTIFY_ONCE` if you want the application to be notified only the first time a procedure executes. Use `NOTIFY_ALWAYS` if you want the application to be notified every time the procedure executes.

Whenever an event with the designated *proc_name* occurs on the Open Server, jConnect calls `eventHdlr.event()` from a separate thread. The event parameters are passed to `eventHdlr.event()` when it is executed. Because it is a separate thread, event notification does not block execution of the application.

If *proc_name* is not a registered procedure, or if Open Server cannot add the client to the notification list, the call to `regWatch()` throws a SQL exception.

To turn off event notification, use this call:

```
SybConnection.regNoWatch(proc_name)
```

Note When you use Sybase event notification extensions, the application needs to call the `close()` method on the connection to remove a child thread created by the first call to `regWatch()`. Failing to do so may cause the Virtual Machine to hang when exiting the application.

Event notification example

The following example shows how to implement an event handler and then register an event with an instance of your event handler, once you have a connection:

```
public class MyEventHandler implements SybEventHandler
{
    // Declare fields and constructors, as needed.
    ...
    public MyEventHandler(String eventname)
    {
        ...
    }

    // Implement SybEventHandler.event.
    public void event(String eventName, ResultSet params)
    {
        try
        {
            // Check for error messages received prior to event
            // notification.
            SQLWarning sqlw = params.getWarnings();
            if sqlw != null
            {
                // process errors, if any
                ...
            }
            // process params as you would any result set with
            // one row.
            ResultSetMetaData rsmd = params.getMetaData();
            int numColumns = rsmd.getColumnCount();
            while (params.next())           // optional
            {
                for (int i = 1; i <= numColumns; i++)
                {
                    System.out.println(rsmd.getColumnName(i) + " =
```



```
        " + params.getString(i));
    }
    // Take appropriate action on the event. For example,
    // perhaps notify application thread.
    ...
}
}
catch (SQLException sqe)
{
    // process errors, if any
    ...
}
}
}

public class MyProgram
{
    ...
    // Get a connection and register an event with an instance
    // of MyEventHandler.
    Connection conn = DriverManager.getConnection(...);
    MyEventHandler myHdlr = new MyEventHandler("MY_EVENT");

    // Register your event handler.
    ((SybConnection)conn).regWatch("MY_EVENT", myHdlr,
        SybEventHandler.NOTIFY_ALWAYS);
    ...
    conn.regNoWatch("MY_EVENT");
    conn.close();
}
```

Handling error messages

jConnect provides two classes for returning Sybase-specific error information, `SySQLException` and `SybSQLWarning`, as well as a `SybMessageHandler` interface that allows you to customize the way jConnect handles error messages received from the server.

Retrieving Sybase-specific error information

jConnect provides an EedInfo interface that specifies methods for obtaining Sybase-specific error information. The EedInfo interface is implemented in SybSQLException and SybSQLWarning, which extend the SQLException and SQLWarning classes.

SybSQLException and SybSQLWarning contain the following methods:

- `public ResultSet getEedParams();`
Returns a one-row result set containing any parameter values that accompany the error message.
- `public int getStatus();`
Returns a “1” if there are parameter values, returns a “0” if there are no parameter values in the message.
- `public int getLineNumber();`
Returns the line number of the stored procedure or query that caused the error message.
- `public String getProcedureName();`
Returns the name of the procedure that caused the error message.
- `public String getServerName();`
Returns the name of the server that generated the message.
- `public int getSeverity();`
Returns the severity of the error message.
- `public int getState();`
Returns information about the internal source of the error message in the server. For use by Sybase Technical Support only.
- `public int getTranState();`
Returns one of the following transaction states:
 - 0 The connection is currently in an extended transaction.
 - 1 The previous transaction committed successfully.
 - 3 The previous transaction aborted.

Some error messages may be `SQLException` or `SQLWarning` messages, without being `SybSQLException` or `SybSQLWarning` messages. Your application should check the type of exception it is handling before it downcasts to `SybSQLException` or `SybSQLWarning`.

Customizing error-message handling

You can use the `SybMessageHandler` interface to customize the way `jConnect` handles error messages generated by the server. Implementing `SybMessageHandler` in your own class for handling error messages can provide the following benefits:

- “Universal” error handling
Error-handling logic can be placed in your error-message handler, instead of being repeated throughout your application.
- “Universal” error logging
Your error-message handler can contain the logic for handling all error logging.
- Remapping of error-message severity, based on application requirements.
Your error-message handler can contain logic for recognizing specific error messages and downgrading or upgrading their severity based on application considerations rather than the server’s severity rating. For example, during a cleanup operation that deletes old rows, you might want to downgrade the severity of a message that a row does not exist; you may want to upgrade the severity in other circumstances.

Note Error-message handlers implementing the `SybMessageHandler` interface only receive server-generated messages. They do not handle messages generated by `jConnect`.

When `jConnect` receives an error message, it checks to see if a `SybMessageHandler` class has been registered for handling the message. If so, `jConnect` invokes the `messageHandler()` method. The `messageHandler()` method accepts a `SQL` exception as its argument, and `jConnect` processes the message based on what value is returned from `messageHandler()`. The error-message handler can:

- Return the `SQL` exception as is.
- Return a null. As a result, `jConnect` ignores the message.

- Create a SQL warning from a SQL exception, and return it. This results in the warning being added to the warning-message chain.
- If the originating message is a SQL warning, `messageHandler()` can evaluate the SQL warning as urgent and create and return a SQL exception to be thrown once control is returned to `jConnect`.

Installing an error-message handler

You can install an error-message handler implementing `SybMessageHandler` by calling the `setMessageHandler()` method from `SybDriver`, `SybConnection`, or `SybStatement`. If you install an error-message handler from `SybDriver`, all subsequent `SybConnection` objects inherit it. If you install an error-message handler from a `SybConnection` object, it is inherited by all `SybStatement` objects created by that `SybConnection`.

This hierarchy only applies from the time the error-message handler object is installed. For example, if you create a `SybConnection` object, *myConnection*, and then call `SybDriver.setMessageHandler()` to install an error-message handler object, *myConnection* cannot use that object.

To return the current error-message handler object, use `getMessageHandler()`.

Error-message-handler example

The following example uses `jConnect` version 5.2.

```
import java.io.*;
import java.sql.*;
import com.sybase.jdbcx.SybMessageHandler;
import com.sybase.jdbcx.SybConnection;
import com.sybase.jdbcx.SybStatement;
import java.util.*;

public class MyApp
{
    static SybConnection conn = null;
    static SybStatement stmt = null
    static ResultSet rs = null;
    static String user = "guest";
    static String password = "sybase";
    static String server = "jdbc:sybase:Tds:192.138.151.39:4444";
    static final int AVOID_SQLE = 20001;

    public MyApp()
```

```

{
    try
    {
        Class.forName("com.sybase.jdbc2.jdbc.SybDriver").newInstance();
        Properties props = new Properties();
        props.put("user", user);
        props.put("password", password);
        conn = (SybConnection)
            DriverManager.getConnection(server, props);
        conn.setMessageHandler(new NoResultSetHandler());
        stmt =(SybStatement) conn.createStatement();
        stmt.executeUpdate("raiserror 20001 'your error'");

        for (SQLWarning sqw = _stmt.getWarnings();
            sqw != null;
            sqw = sqw.getNextWarning());
        {
            if (sqw.getErrorCode() == AVOID_SQLLE);
            {
                System.out.println("Error" + sqw.getErrorCode()+
                    " was found in the Statement's warning list.");
                break;
            }
        }
        stmt.close();
        conn.close();
    }
    catch(Exception e)
    {
        System.out.println(e.getMessage());
        e.printStackTrace();
    }
}

class NoResultSetHandler implements SybMessageHandler
{
    public SQLException messageHandler(SQLException sqe)
    {
        int code = sqe.getErrorCode();
        if (code == AVOID_SQLLE)
        {
            System.out.println("User " + _user + " downgrading " +
                AVOID_SQLLE + " to a warning");
            sqe = new SQLWarning(sqe.getMessage(),
                sqe.getSQLState(),sqe.getErrorCode());
        }
    }
}

```

```
        return sqe;
    }
}

public static void main(String args[])
{
    new MyApp();
}
```

Storing Java objects as column data in a table

Some database products enable you to directly store Java objects as column data in a database. In such databases, Java classes are treated as datatypes, and you can declare a column with a Java class as its datatype.

jConnect supports storing Java objects in a database by implementing the `setObject()` methods defined in the `PreparedStatement` interface and the `getObject()` methods defined in the `CallableStatement` and `ResultSet` interfaces. This allows you to use jConnect with an application that uses native JDBC classes and methods to directly store and retrieve Java objects as column data.

Note To use `getObject()` and `setObject()`, set the jConnect version to `VERSION_4` or later. See “Setting the jConnect version” on page 6.

The following sections describe the requirements and procedures for storing objects in a table and retrieving them using JDBC with jConnect:

- Prerequisites for storing java objects as column data
- Sending Java objects to a database
- Receiving Java objects from the database

Note Adaptive Server version 12.0 and later and Adaptive Server Anywhere version 6.0.x and later can store Java objects in a table, with some limitations. See the *jConnect for JDBC Release Bulletin* for more information.

Prerequisites for storing java objects as column data

To store Java objects belonging to a user-defined Java class in a column, three requirements must be met:

- The class must implement the `java.io.Serializable` interface. This is because `jConnect` uses native Java serialization and deserialization to send objects to a database and receive them back from the database.
- The class definition must be installed in the destination database, *or* you must be using the `DynamicClassLoader` (DCL) to load a class directly from an Adaptive Server Anywhere or an Adaptive Server server and use it as if it was present in the local `CLASSPATH`. See “Dynamic class loading” on page 80 for more information.
- The client system must have the class definition in a `.class` file that is accessible through the local `CLASSPATH` environment variable.

Sending Java objects to a database

To send an instance of a user-defined class as column data, use one of the following `setObject()` methods, as specified in the `PreparedStatement` interface:

```
void setObject(int parameterIndex, Object x, int targetSqlType,
    int scale) throws SQLException;
void setObject(int parameterIndex, Object x, int targetSqlType)
    throws SQLException;
void setObject(int parameterIndex, Object x) throws SQLException;
```

In `jConnect 5.5`, to send a Java object, you can use the `java.sql.Types.JAVA_OBJECT` target `sqlType`, or you can use `java.sql.Types.OTHER`.

The following example defines an `Address` class, shows the definition of a `Friends` table that has an `Address` column whose datatype is the `Address` class, and inserts a row into the table.

```
public class Address implements Serializable
{
    public String    streetNumber;
    public String    street;
    public String    apartmentNumber;
    public String    city;
    public int       zipCode;
    //Methods
    ...
}

/* This code assumes a table with the following structure
** Create table Friends:
** (firstname varchar(30) ,
**  lastname varchar(30),
```

```
**  address Address,
**  phone varchar(15))
*/

// Connect to the database containing the Friends table.
Connection conn =
    DriverManager.getConnection("jdbc:sybase:Tds:localhost:5000",
        "username", "password");

// Create a Prepared Statement object with an insert statement
//for updating the Friends table.
PreparedStatement ps = conn.prepareStatement("INSERT INTO
    Friends values (?, ?, ?, ?)");

// Now, set the values in the prepared statement object, ps.
// set firstname to "Joan."
ps.setString(1, "Joan");

// Set last name to "Smith."
ps.setString(2, "Smith");

// Assuming that we already have "Joan_address" as an instance
// of Address, use setObject(int parameterIndex, Object x) to
// set the address column to "Joan_address."
ps.setObject(3, Joan_address);

// Set the phone column to Joan's phone number.
ps.setString(4, "123-456-7890");

// Perform the insert.
ps.executeUpdate();
```

Receiving Java objects from the database

A client JDBC application can receive a Java object from the database in a result set or as the value of an output parameter returned from a stored procedure.

If a result set contains a Java object as column data, use one of the following `getObject()` methods in the `ResultSet` interface to retrieve the object:

```
Object getObject(int columnIndex) throws SQLException;
Object getObject(String columnName) throws SQLException;
```


If an output parameter from a stored procedure contains a Java object, use the following `getObject()` method in the `CallableStatement` interface to retrieve the object:

`Object getObject(int parameterIndex)` throws `SQLException`;

The following example illustrates the use of `ResultSet.getObject(int parameterIndex)` to assign an object received in a result set to a class variable. The example uses the `Address` class and `Friends` table used in the previous section and presents a simple application that prints a name and address on an envelope.

```
/*
** This application takes a first and last name, gets the
** specified person's address from the Friends table in the
** database, and addresses an envelope using the name and
** retrieved address.
*/
public class Envelope
{
    Connection conn = null;
    String firstName = null;
    String lastName = null;
    String street = null;
    String city = null;
    String zip = null;

    public static void main(String[] args)
    {
        if (args.length < 2)
        {
            System.out.println("Usage: Envelope <firstName>
                               <lastName>");
            System.exit(1);
        }
        // create a 4" x 10" envelope
        Envelope e = new Envelope(4, 10);
        try
        {
            // connect to the database with the Friends table.
            conn = DriverManager.getConnection(
                "jdbc:sybase:Tds:localhost:5000", "username",
                "password");
            // look up the address of the specified person
            firstName = args[0];
            lastName = args[1];
            PreparedStatement ps = conn.prepareStatement(
```

```
        "SELECT address FROM friends WHERE " +
        "firstname = ? AND lastname = ?");
    ps.setString(1, firstName);
    ps.setString(2, lastName);
    ResultSet rs = ps.executeQuery();
    if (rs.next())
    {
        Address a = (Address) rs.getObject(1);
        // set the destination address on the envelope
        e.setAddress(firstName, lastName, a);
    }
    conn.close();
}
catch (SQLException sqe)
{
    sqe.printStackTrace();
    System.exit(2);
}
// if everything was successful, print the envelope
e.print();
}
private void setAddress(String fname, String lname, Address a)
{
    street = a.streetNumber + " " + a.street + " " +
        a.apartmentNumber;
    city = a.city;
    zip = " " + a.zipCode;
}
private void print()
{
    // Print the name and address on the envelope.
    ...
}
}
```

You can find a more detailed example of `HandleObject.java` in the *sample2* subdirectory under your `jConnect` installation directory.

Dynamic class loading

Adaptive Server Anywhere version 6.0 and Adaptive Server version 12.0 and later allow you to specify Java classes as:

- Datatypes of SQL columns

- Datatypes of Transact-SQL variables
- Default values for SQL columns

In earlier versions, only classes that appeared in jConnect's CLASSPATH were accessible; that is, if a jConnect application attempted to access an instance of a class that was not in the local CLASSPATH, a `java.lang.ClassNotFoundException` exception would result.

jConnect version 5.x and later implements `DynamicClassLoader` (DCL) to load a class directly from an Adaptive Server Anywhere or Adaptive Server server and use it as if it was present in the local CLASSPATH.

All security features present in the superclass are inherited. The loader delegation model implemented in Java 2 is followed—first jConnect attempts to load a requested class from the CLASSPATH; if that fails, jConnect tries the `DynamicClassLoader`.

See *Java in Adaptive Server* for more detailed information about use Java and Adaptive Server.

Using *DynamicClassLoader*

To use DCL functionality:

- 1 Create and configure a class loader. Your jConnect application's code should look similar to this:

```
Properties props = new Properties();

// URL of the server where the classes live.
String classesUrl = "jdbc:sybase:Tds:myase:1200";

// Connection properties for connecting to above server.
props.put("user", "grinch");
props.put("password", "meanone");
...

// Ask the SybDriver for a new class loader.
DynamicClassLoader loader = driver.getClassLoader(classesUrl, props);
```

- 2 Use the `CLASS_LOADER` connection property to make the new class loader available to the statement that executes the query. Once you create the class loader, pass it to subsequent connections as shown below (continuing from the code example in step 1).

```
// Stash the class loader so that other connection(s)
// can know about it.
props.put("CLASS_LOADER", loader);
```

```
// Additional connection properties
props.put("user", "joeuser");
props.put("password", "joespassword");

// URL of the server we now want to connect to.
String url = "jdbc:sybase:Tds:jdbc.sybase.com:4446";

// Make a connection and go.
Connection conn = DriverManager.getConnection(url, props);
```

Assuming the following Java class definition:

```
class Addr {
    String street;
    String city;
    String state;
}
```

and the following SQL table definition:

```
create table employee (char(100) name, int empid, Addr address)
```

- 3 Use the following client-side code in the absence of an Addr class in the client application's CLASSPATH:

```
Statement stmt = conn.createStatement();
// Retrieve some rows from the table that has a Java class
// as one of its fields.
ResultSet rs = stmt.executeQuery(
    "select * from employee where empid = '19'");
if (rs.next()) {
    // Even though the class is not in our class path,
    // we should be able to access its instance.
    Object obj = rs.getObject("address");
    // The class has been loaded from the server,
    // so let's take a look.
    Class c = obj.getClass();

    // Some Java Reflection can be done here
    // to access the fields of obj.
    ...
}
```

The `CLASS_LOADER` connection property provides a convenient mechanism for sharing one class loader among several connections.

You should ensure that sharing a class loader across connections does not result in class conflicts. For example, if two different, incompatible instances of class `org.foo.Bar` exist in two different databases, problems can arise if you use the same loader to access both classes. The first class is loaded when examining a result set from the first connection. When it is time to examine a result set from the second connection, the class is already loaded. The second class is never loaded, and there is no direct way for `jdbc:connect` to detect this situation.

However, Java has a built-in mechanism for ensuring that the version of a class matches the version information in a deserialized object. The above situation is at least detected and reported by Java.

Classes and their instances do not need to reside in the same database or server, but there is no reason why both the loader and subsequent connections cannot refer to the same database/server.

Deserialization

The following example illustrates how to deserialize an object from a local file. The serialized object is an instance of a class that resides on a server and does not exist in the `CLASSPATH`.

`SybResultSet.getObject()` makes use of `DynamicObjectInputStream`, which is a subclass of `ObjectInputStream` that loads a class definition from `DynamicClassLoader`, rather than the default system (“boot”) class loader.

```
// Make a stream on the file containing the
//serialized object.
FileInputStream fileStream = new FileInputStream("serFile");
// Make a "deserializer" on it. Notice that, apart
//from the additional parameter, this is the same
//as ObjectInputStreamDynamicObjectInputStream
stream = new DynamicObjectInputStream(fileStream, loader);
// As the object is deserialized, its class is
//retrieved through the loader from our server.
Object obj = stream.readObject();stream.close();
```

Preloading JARS

`jdbc:connect` version 5.x and later have a connection property called `PRELOAD_JARS`. When defined as a comma-delimited list of JAR file names, the JAR files are loaded in their entirety. In this context, “JAR” refers to the “retained JARname” used by the server. This is the JAR name specified in the install Java program, for example:

```
install java new jar 'myJarName' from file '/tmp/mystuff.jar'
```

If you set `PRELOAD_JARS`, the JAR files are associated with the class loader, so it is unnecessary to preload them with every connection. You should only specify `PRELOAD_JARS` for one connection. Subsequent attempts to preload the same JAR files may result in performance problems as the JAR data is retrieved from the server unnecessarily.

Note Adaptive Server Anywhere 6.x and later cannot return a JAR file as one entity, so `jConnect` iteratively retrieves each class in turn. However, Adaptive Server 12.x and later retrieves the entire JAR and loads each class that it contains.

Advanced features

There are various public methods in `DynamicClassLoader`. For more information, see the javadocs information in *JDBC_HOME/docs/en/javadocs*.

Additional features include the ability to keep a loader's database connection "alive" when a series of class loads is expected, and to explicitly load a single class by name.

Public methods inherited from `java.lang.ClassLoader` can also be used. Methods in `java.lang.Class` that deal with loading classes are also available; however, use these methods with caution since some of them make assumptions about which class loader gets used. In particular, you should use the 3-argument version of `Class.forName()`, otherwise the system ("boot") class loader will be used. See "Handling error messages" on page 71.

JDBC 2.0 optional package extensions support

The JDBC 2.0 Optional Package (formerly the JDBC 2.0 Standard Extension API) defined several features that JDBC 2.0 drivers could implement. `jConnect` version 5.x and later have implemented the following optional package extension features:

- JNDI for naming databases
(works with any Sybase DBMS supported by `jConnect`)
- Connection pooling
(works with any Sybase DBMS supported by `jConnect`)

- Distributed transaction management support (works only with Adaptive Server version 12.0 and later, or version 11.x using XA-Server™)

The above features require classes and/or interfaces that are not found in standard JDK 1.2.x distributions. You must download `javax.sql.*` and `javax.naming.*` to implement them if you're using a JDK 1.2.x or JRE installation. However, if you are using JDK 1.3.x or later, no additional download is necessary since the classes are part of the standard Java installation Databases and Connection Pooling, and you must download `javax.transaction.xa.*` to implement Distributed Transaction Management Support.

Note Sybase recommends that you use JNDI 1.2, which is compatible with Java 1.1.6 and later.

JNDI for naming databases

Reference

The *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*), Chapter 5, “JNDI and the JDBC API.”

Related interfaces

- `javax.sql.DataSource`
- `javax.naming.Referenceable`
- `javax.naming.spi.ObjectFactory`

This feature provides JDBC clients with an alternative to the standard approach for obtaining database connections. Instead of invoking `Class.forName(“com.sybase.jdbc2.jdbc.SybDriver”)`, then passing a JDBC URL to the `DriverManager`'s `getConnection()` method, clients can access a JNDI name server using a logical name to retrieve a `javax.sql.DataSource` object. This object is responsible for loading the driver and establishing the connection to the physical database it represents. The client code is simpler and reusable because the vendor-specific information has been placed within the `DataSource` object.

The Sybase implementation of the `DataSource` object is `com.sybase.jdbcx.SybDataSource` (see the javadocs for details). This implementation supports the following standard properties using the design pattern for JavaBean components:

- `databaseName`
- `dataSourceName`
- `description`
- `networkProtocol`
- `password`
- `portNumber`
- `serverName`
- `user`

`roleName` is not supported.

`jConnect` provides an implementation of the `javax.naming.spi.ObjectFactory` interface so the `DataSource` object can be constructed from the attributes of a name server entry. When given a `javax.naming.Reference`, or a `javax.naming.Name` and a `javax.naming.DirContext`, this factory can construct `com.sybase.jdbcx.SybDataSource` objects. To use this factory, set the `java.naming.object.factory` system property to include `com.sybase.jdbc2.SybObjectFactory`.

Usage

You can use `DataSource` in different ways, in different applications. All options are discussed below with some code examples to guide you through the process. For more information, see the *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*), and the JNDI documentation on Sun's Web site.

1a. Configuration by administrator: LDAP

`jConnect` has supported LDAP connectivity since version 4.0. As a result, the recommended approach, which requires no custom software, is to configure `DataSources` as LDAP entries using the LDAP Data Interchange Format (LDIF). For example:

```
dn:servername:myASE, o=MyCompany, c=US
1.3.6.1.4.1.897.4.2.5:TCP#1# mymachine 4000
1.3.6.1.4.1.897.4.2.10:PACKETSIZE=1024&user=me&password=secret
1.3.6.1.4.1.897.4.2.11:userdb
```


1b. Access by client

This is the typical JDBC client application. The only difference is that you access the name server to obtain a reference to a `DataSource` object, instead of accessing the `DriverManager` and providing a JDBC URL. Once you obtain the connection, the client code is identical to any other JDBC client code. The code is very generic and references Sybase only when setting the object factory property, which can be set as part of the environment.

The jConnect installation contains the sample program *sample2/SimpleDataSource.java* to illustrate the use of `DataSource`. This sample is provided for reference only; that is, you cannot run the sample unless you configure your environment and edit the sample appropriately. *SimpleDataSource.java* contains the following critical code:

```
import javax.naming.*;
import javax.sql.*;
import java.sql.*;

// set necessary JNDI properties for your environment (same as above)
Properties jndiProps = new Properties();

// used by JNDI to build the SybDataSource
jndiProps.put(Context.OBJECT_FACTORIES,
    "com.sybase.jdbc2.jdbc.SybObjectFactory");

// nameserver that JNDI should talk to
jndiProps.put(Context.PROVIDER_URL, "ldap:
//some_ldap_server:238/o=MyCompany,c=Us");

// used by JNDI to establish the naming context
jndiProps.put(Context.INITIAL_CONTEXT_FACTORY,
    "com.sun.jndi.ldap.LdapCtxFactory");

// obtain a connection to your name server
Context ctx = new InitialContext(jndiProps);
DataSource ds = (DataSource) ctx.lookup("servername=myASE");

// obtains a connection to the server as configured earlier.
// in this case, the default username and password will be used
Connection conn = ds.getConnection();

// do standard JDBC methods
...
```

Explicitly passing the `Properties` to the `InitialContext` constructor is not required if the properties have already been defined within the virtual machine; that is, passed when Java was either set as part of the browser properties, or by using:

```
java -Djava.naming.object.factory=com.sybase.jdbc2.jdbc.SybObjectFactory
```

See your Java VM documentation for more information about setting environment properties.

2a. Configuration by administrator: custom

This phase is typically done by the person who does database system administration or application integration for their company. The purpose is to define a data source, then deploy it under a logical name to a name server. If the server needs to be reconfigured (for example, moved to another machine, port, and so on), then the administrator runs this configuration utility (outlined below) and reassigns the logical name to the new data source configuration. As a result, the client code does not change, since it knows only the logical name.

```
import javax.sql.*;
import com.sybase.jdbcx.*;
.....

// create a SybDataSource, and configure it
SybDataSource ds = new com.sybase.jdbc2.jdbc.SybDataSource();
ds.setUser("my_username");
ds.setPassword("my_password");
ds.setDatabaseName("my_favorite_db");
ds.setServerName("db_machine");
ds.setPortNumber(4000);
ds.setDescription("This DataSource represents the Adaptive Server
    Enterprise server running on db_machine at port 2638. The default
    username and password have been set to 'me' and 'mine' respectively.
    Upon connection, the user will access the my_favorite_db database on
    this server.");
Properties props = new Properties();
props.put("REPEAT_READ", "false");
props.put("REQUEST_HA_SESSION", "true");
ds.setConnectionProperties(props);
// store the DataSource object. Typically this is
// done by setting JNDI properties specific to the
// type of JNDI service provider you are using.
// Then, initialize the context and bind the object.
Context ctx = new InitialContext();
ctx.bind("jdbc/myASE", ds);
```

Once you set up your `DataSource`, you decide where and how you want to store the information. To assist you, `SybDataSource` is both `java.io.Serializable` and `javax.naming.Referenceable`, but it is still up to the administrator to determine how the data is stored depending on what service provider you are using for JNDI.

2b. Access by client

The client retrieves the `DataSource` object by setting its JNDI properties the same way the `DataSource` was deployed. The client needs to have an object factory available that can transform the object as it is stored (for example, serialized) into a Java object.

```
Context ctx = new InitialContext();
DataSource ds = (DataSource ctx.lookup("jdbc/myASE"));
```

Connection pooling

Reference

The *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*), Chapter 6, “Connection Pooling.”

Related interfaces

- `javax.sql.ConnectionPoolDataSource`
- `javax.sql.PooledConnection`

Overview

Traditional database applications create one connection to a database that you use for each session of an application. However, a Web-based database application may need to open and close a new connection several times during the application’s use. An efficient way to handle Web-based database connections is to use connection pooling, which maintains open database connections and manages connection sharing across different user requests to maintain performance and to reduce the number of idle connections. On each connection request, the connection pool first determines if there is an idle connection in the pool. If there is, the connection pool returns that connection instead of making a new connection to the database.

Connection pooling capabilities are provided by `ConnectionPoolDataSource`. If you use this interface, you can pool connections. If you use the `DataSource` interface, you cannot pool connections.

When you use `ConnectionPoolDataSource`, pool implementations listen to the `PooledConnection`. The implementation is notified when a user closes the connection, or if the user has an error that destroys the connection. At this point, the pool implementation decides what to do with the `PooledConnection`.

Without connection pooling, a transaction:

- 1 Creates a connection to the database.

- 2 Sends the query to the database.
- 3 Gets back the result set.
- 4 Displays the result set.
- 5 Destroys the connection.

With connection pooling, the sequence looks more like this:

- 1 Sees if an unused connection exists in the “pool” of connections.
- 2 If so, uses it; otherwise creates a new connection.
- 3 Sends the query to the database.
- 4 Gets back the result set.
- 5 Displays the result set.
- 6 Returns the connection to the “pool.” The user still calls “close()”, but the connection remains open and the pool is notified of the close request.

It is less costly to reuse a connection than to create a new one every time a client needs to establish a connection to a database.

To enable a third party to implement the connection pool, the `jConnect` implementation has the `ConnectionPoolDataSource` interface produce `PooledConnections`, similar to how the `DataSource` interface produces `Connections`.

The pool implementation creates “real” database connections, using the `getPooledConnection()` methods of `ConnectionPoolDataSource`. Then, the pool implementation registers itself as a listener to the `PooledConnection`.

Currently, when a client requests a connection, the pool implementation invokes `getConnection()` on an available `PooledConnection`. When the client finishes with the connection and calls `close()`, the pool implementation is notified through the `ConnectionEventListener` interface that the connection is free and available for reuse.

The pool implementation is also notified through the `ConnectionEventListener` interface if the client somehow corrupts the database connection, so that the pool implementation can remove that connection from the pool.

For more information, refer to Appendix B of the *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*).

Configuration by administrator: LDAP

This approach is the same as “1a. Configuration by administrator: LDAP” described in “JNDI for naming databases,” except that you enter an additional line to your LDIF entry. In the following example, the added line of code is bolded for your reference.

```
dn:servername=myASE, o=MyCompany, c=US
1.3.6.1.4.1.897.4.2.5:TCP#1# mymachine 4000
1.3.6.1.4.1.897.4.2.10:PACKETSIZE=1024&user=me&password=secret
1.3.6.1.4.1.897.4.2.11:userdb
1.3.6.1.4.1.897.4.2.18:ConnectionPoolDataSource
```

Access by middle-tier clients

This procedure initializes three properties (INITIAL_CONTEXT_FACTORY, PROVIDER_URL, and OBJECT_FACTORIES as shown on page 78), and retrieves a ConnectionPoolDataSource object. For a more complete code example, see *sample2/SimpleConnectionPool.java*. The fundamental difference is:

```
...
ConnectionPoolDatabase cpds = (ConnectionPoolDataSource)
    ctx.lookup("servername=myASE");
PooledConnection pconn = cpds.getPooledConnection();
```

Distributed transaction management support

This feature provides a standard Java API for performing distributed transactions with either Adaptive Server version 12.x or version 11.x with XA-Server.

Note This feature is designed for use in a large multitier environment.

Reference

See Chapter 7, “Distributed Transactions,” in the *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*).

Related interfaces


- `javax.sql.XADataSource`
- `javax.sql.XAConnection`
- `javax.transaction.xa.XAResource`

Background and system requirements

For Adaptive Server
12.0 and later

- Because jConnect is communicating directly with the resource manager within Sybase Adaptive Server version 12.0 and later, the installation must have Distributed Transaction Management support.
- Any user that wants to participate in a distributed transaction must have the “dtm_tm_role” granted to them or the transactions will fail.
- To use distributed transactions, you must install the stored procedures in the `/sp` directory. Refer to “Installing Stored Procedures” in Chapter 1 of your *jConnect for JDBC Installation Guide*.


Figure 2-2: Distributed transaction management support with version 12.x



For Adaptive Server
11.x

jConnect also provides a standard Java API for performing distributed transactions with Adaptive Server version 11.x as your database server.

- This implementation works only with Sybase Adaptive Server version 11.x and XA-Server 11.1.

Figure 2-3: Distributed transaction management support with version 11.x

- The login chosen cannot have a default login database of master, model, or sybsystemdb. This is because XA-Server connects only when the user's work is associated with a distributed transaction, and distributed transactions are not permitted on those databases.
- There is no access to metadata. While this restricts the client, it is most likely not the part of the API being used within the boundaries of distributed transactions.

Adaptive Server 12.x use

Configuration by
administrator: LDAP

This approach is the same as “1a. Configuration by administrator: LDAP” described in “JNDI for naming databases” on page 85, except that you enter an additional line to the LDIF entry. In the following example, the added line of code is displayed in bold.

```

dn:servername:myASE, o=MyCompany, c=US
1.3.6.1.4.1.897.4.2.5:TCP#1# mymachine 4000
1.3.6.1.4.1.897.4.2.10:PACKETSIZE=1024&user=me&password=secret
1.3.6.1.4.1.897.4.2.11:userdb
1.3.6.1.4.1.897.4.2.18:XADataSource
  
```

Access by middle-tier
clients

This procedure initializes three properties (INITIAL_CONTEXT_FACTORY, PROVIDER_URL, and OBJECT_FACTORIES), and retrieves a XADataSource object. For example:

```

...
XADataSource xads = (XADataSource) ctx.lookup("servername=myASE");
XAConnection xaconn = xads.getXAConnection();
  
```

or override the default settings for the user name and password:

```

...
XADataSource xads = (XADataSource) ctx.lookup("servername=myASE");
  
```

```
XAConnection xaconn = xads.getXAConnection("my_username","my_password");
```

Adaptive Server 11.x use

Configuration by
administrator: LDAP

This approach is the same as “1a. Configuration by administrator: LDAP” described in “JNDI for naming databases” on page 85, except that you enter an additional three lines to the LDIF entry.

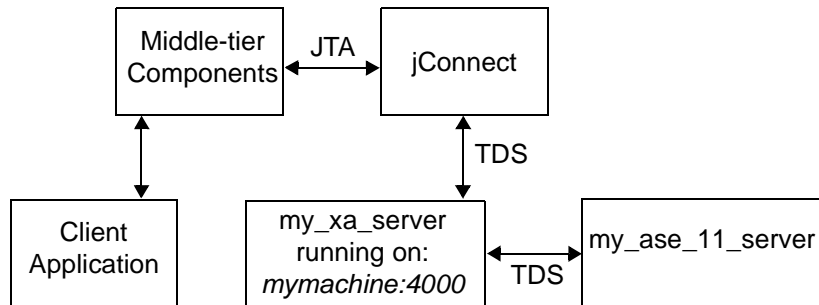
In the following example, the additional code lines are displayed in bold.

```
dn:servername:myASE, o=MyCompany, c=US
1.3.6.1.4.1.897.4.2.5:TCP#1# mymachine 4000
1.3.6.1.4.1.897.4.2.10:PACKETSIZE=1024&user=me&password=secret
1.3.6.1.4.1.897.4.2.11:userdb
1.3.6.1.4.1.897.4.2.16:userconnection
1.3.6.1.4.1.897.4.2.17:1
1.3.6.1.4.1.897.4.2.18:XADatasource
```

where . . .4.2.17:1 indicates that jConnect is going to connect to an XA-Server and *userconnection* corresponds to the Logical Resource Manager (LRM) to use. XA-Server has an *xa_config* file that contains these entries:

```
[xa]
lrn=userconnection
server=my_ase_11_server
XAserver=my_xa_server
```

Figure 2-4: Distributed transaction management support sample configuration



See the XA-Server documentation for details on how to write an *xa_config* file.

Access by middle-tier
clients

This procedure initializes three properties (INITIAL_CONTEXT_FACTORY, PROVIDER_URL, and OBJECT_FACTORIES), and retrieves a XADatasource object. For example:

```
...
XADatasource xads = (XADatasource) ctx.lookup("servername=myASE");
```



```
XAConnection xaconn = xads.getXAConnection();
```

With Adaptive Server 11.x, you *cannot* override the default user name and password; that is, you cannot call:

```
xads.getXAConnection("my_username", "my_password");
```

because the *lrm* is associated with a specific user name and password.

Restrictions on and interpretations of JDBC standards

This section discusses restrictions on and interpretations of JDBC standards that apply to jConnect, including how the jConnect implementation of JDBC deviates from the JDBC 1.x and 2.0 standards. The following topics are covered:

- `Connection.isClosed()` and the `IS_CLOSED_TEST` connection property
- `Statement.close()` with unprocessed results
- Making adjustments for multithreading
- Using `ResultSet.getCursorName()`
- Using `setLong()` with large parameter values
- Using `COMPUTE` statements
- Executing stored procedures

Connection.isClosed() and the IS_CLOSED_TEST connection property

According to section 11.1 of the JDBC 2.1 specification:

“The `Connection.isClosed()` method is only guaranteed to return true after `Connection.close()` has been called. `Connection.isClosed()` cannot be called, in general, to determine if a database connection is valid or invalid. A typical client can determine that a connection is invalid by catching the exception that is thrown when an operation is attempted.”

jConnect offers a default interpretation of the `isClosed()` method that is different from the behavior that is defined in the spec. When you call `Connection.isClosed()`, jConnect will first check if `Connection.close()` has been called on this connection. If `close()` has been called, jConnect will return `true` for `isClosed()`.

However, if `Connection.close()` has not been called, jConnect will next try to execute the `sp_mda` stored procedure on the database. The `sp_mda` stored procedure is part of the standard metadata that jConnect users must install when they use jConnect with a database.

The purpose of calling `sp_mda` is so that jConnect can try to execute a procedure that is known (or at least, expected) to reside on the database server. If the stored procedure executes normally, then jConnect returns *false* for `isClosed()` because we have verified that the database `Connection` is valid and working. However, if the call to `sp_mda` results in a `SQLException` being thrown, jConnect catches the exception and returns `true` for `isClosed()` because it appears that there is something wrong with the `Connection`.

If you wish to force jConnect to more closely follow the standard JDBC behavior for `isClosed()`, you can do so by setting the `IS_CLOSED_TEST` connection property to the special value `'INTERNAL'`. The `INTERNAL` setting means that jConnect will return `true` for `isClosed()` only when `Connection.close()` has been called, or when jConnect has detected an `IOException` that has disabled the `Connection`.

You can also specify a query other than `sp_mda` to use when `isClosed()` is called. For example, if you wanted jConnect to try a *select 1* when `isClosed()` was called, you could set the `IS_CLOSED_TEST` connection property to *select 1*.

Statement.close() with unprocessed results

The JDBC specification is somewhat vague on how a driver should behave when you call `Statement.execute()` and later call `close()` on that same `Statement` object without processing all of the results (update counts and `ResultSet`s) returned by the `Statement`.

For example, assume that there is a stored procedure on the database that does seven row inserts. An application then executes that stored procedure using a `Statement.execute()`. In this case, a Sybase database will return seven update counts (one for each inserted row) to the application. In normal JDBC application logic, you would process those update counts in a loop using the `getMoreResults()`, `getResultSet()` and `getUpdateCount()` methods. These are clearly explained on the *java.sun.com* website in the javadocs for the `java.sql.*` package.

An application programmer, however, might incorrectly choose to call `Statement.close()` before reading through all of the returned update counts. In this case, `jdbcConnect` will send a cancel to the database, which could have unexpected and unwanted side effects.

In this particular example, if the application called `Statement.close()` before the database had completed the inserts, the database might not execute all of the inserts. It might stop, for example, after only five rows were inserted because the cancel would be processed on the database before the stored procedure completed.

The missing inserts would not be reported to you in this case. Future releases of `jdbcConnect` may throw a `SQLException` when you try to close a `Statement` when there are still unprocessed results, but until then, `jdbcConnect` programmers are strongly advised to adhere to the following guidelines:

- When you call `Statement.close()`, a cancel is sent to the server if not all the results (update counts and `ResultSets`) have been completely processed by you. In cases where you only executed select statements, this is fine. However, in cases where you executed insert/update/delete operations, this could result in not all of those operations completing as expected.
- Therefore, you should never call `close()` with unprocessed results when you have executed anything but pure select statements.
- Instead, if you call `Statement.execute()` be sure your code processes all the results by using the `getUpdateCount()`, `getMoreResults()` and `getResultSet()` methods.

Making adjustments for multithreading

If several threads simultaneously call methods on the same `Statement` instance, `CallableStatement`, or `PreparedStatement`—which we do not recommend—you have to manually synchronize the calls to the methods on the `Statement`; `jdbcConnect` does not do this automatically.

For example, if you have two threads operating on the same `Statement` instance—one thread sending a query and the other thread processing warnings—you have to synchronize the calls to the methods on the `Statement` or conflicts may occur.

Using `ResultSet.setCursorName()`

Some JDBC drivers generate a cursor name for any SQL query so that a string can always be returned. However, `jConnect` does not return a name when `ResultSet.setCursorName()` is called, unless you either

- called `setFetchSize()` or `setCursorName()` on the corresponding `Statement`, or
- set the `SELECT_OPENS_CURSOR` connection property to true, and your query was in the form of `SELECT... FOR UPDATE`; for example,

```
select au_id from authors for update
```

If you do not call `setFetchSize()` or `setCursorName()` on the corresponding `Statement`, or set the `SELECT_OPENS_CURSOR` connection property to true, null is returned.

According to the JDBC 2.0 API (chapter 11, “Clarifications”), all other SQL statements do not need to open a cursor and return a name.

For more information on how to use cursors in `jConnect`, see “Using cursors with result sets” on page 48.

Using `setLong()` with large parameter values

Implementations of the `PreparedStatement.setLong()` method set a parameter value to a SQL `BIGINT` datatype. Most Adaptive Server databases do not have an 8-byte `BIGINT` datatype. If a parameter value requires more than 4 bytes of a `BIGINT`, using `setLong()` may result in an overflow exception.

Using `COMPUTE` statements

`jConnect` does not support computed rows. Results are automatically cancelled when a query contains a computed row. For example, the following statement is rejected:

```
SELECT name FROM sysobjects
WHERE type="S" COMPUTE COUNT(name)
```

To avoid this problem, substitute the following code:

```
SELECT name from sysobjects WHERE type="S"
SELECT COUNT(name) from sysobjects WHERE type="S"
```

Executing stored procedures

- If you execute a stored procedure in a `CallableStatement` object that represents parameter values as question marks, you get better performance than if you use both question marks and literal values for parameters. Further, if you mix literals and question marks, you cannot use output parameters with a stored procedure.

The following example creates `sp_stmt` as a `CallableStatement` object for executing the stored procedure `MyProc`:

```
CallableStatement sp_stmt = conn.prepareCall(
    "{call MyProc(?,?)}");
```

The two parameters in `MyProc` are represented as question marks. You can register one or both of them as output parameters using the `registerOutParameter()` methods in the `CallableStatement` interface.

In the following example, `sp_stmt2` is a `CallableStatement` object for executing the stored procedure `MyProc2`.

```
CallableStatement sp_stmt2 = conn.prepareCall(
    {"call MyProc2(?, 'javelin')"});
```

In `sp_stmt2`, one parameter value is given as a literal value and the other as a question mark. You cannot register either parameter as an output parameter.

- To execute stored procedures with RPC commands using name-binding for parameters, use either of the following procedures.
 - Use language commands, passing input parameters to them directly from Java variables using the `PreparedStatement` class. This is illustrated in the following code fragment:

```
// Prepare the statement
System.out.println("Preparing the statement...");
String stmtString = "exec " + procname + " @p3=?, @p1=?";
PreparedStatement pstmt = con.prepareStatement(stmtString);
```

```
// Set the values
pstmt.setString(1, "xyz");
pstmt.setInt(2, 123);

// Send the query
System.out.println("Executing the query...");
ResultSet rs = pstmt.executeQuery();
```

- With jConnect version 5.x and later, use the `com.sybase.jdbcx.SybCallableStatement` interface, illustrated in this example:

```
import com.sybase.jdbcx.*;
....
// prepare the call for the stored procedure to execute as an RPC
String execRPC = "{call " + procName + " (?, ?)}";
SybCallableStatement scs = (SybCallableStatement)
con.prepareCall(execRPC);

// set the values and name the parameters
// also (optional) register for any output parameters
scs.setString(1, "xyz");
scs.setParameterName(1, "@p3");
scs.setInt(2, 123);
scs.setParameterName(2, "@p1");

// execute the RPC
// may also process the results using getResultSet()
// and getMoreResults()

// see the samples for more information on processing results
ResultSet rs = scs.executeQuery();
```

This chapter describes solutions and workarounds for problems you might encounter when using jConnect.

Topics

Debugging with jConnect

Capturing TDS communication

Unsuccessful connection errors

Memory usage in jConnect applications

Stored procedure errors

Custom socket implementation error

Debugging with jConnect

jConnect includes a `Debug` class that contains a set of debugging functions. The `Debug` methods include a variety of `assert`, `trace`, and `timer` functions that let you define the scope of the debugging process and the output destination for the debugging results.

The jConnect installation also includes a complete set of debug-enabled classes. These classes are located in the *devclasses* subdirectory under your jConnect installation directory. For debugging purposes, you must redirect your `CLASSPATH` environment variable to reference the debug mode runtime classes (*devclasses/jconn2d.jar*), rather than the standard jConnect *classes* directory. You can also do this by explicitly providing a `-classpath` argument to the `java` command when you run a Java program.

Obtaining an instance of the `Debug` class

To use the jConnect debugging feature, your application must import the `Debug` interface and obtain an instance of the `Debug` class by calling the `getDebug()` method on the `SybDriver` class.

```
import com.sybase.jdbcx.Debug
import com.sybase.jdbcx.SybDebug
//
...
SybDriver sybDriver = (SybDriver)

Class.forName("com.sybase.jdbc2.jdbc.SybDriver").newInstance();
Debug sybdebug = sybDriver.getDebug();
...
```

Turning on debugging in your application

To use the `debug()` method on the `Debug` object to turn on debugging within your application, add this call:

```
sybdebug.debug(true, [classes], [printstream]);
```

The *classes* parameter is a string that lists the specific classes you want to debug, separated by colons. For example:

```
sybdebug.debug(true, "MyClass")
```

and

```
sybdebug.debug(true, "MyClass:YourClass")
```

“STATIC” in the class string turns on debugging for all static methods in jConnect in addition to the designated classes. For example:

```
sybdebug.debug(true, "STATIC:MyClass")
```

You can specify “ALL” to turn on debugging for all classes. For example:

```
sybdebug.debug(true, "ALL");
```

The *printstream* parameter is optional. If you do not specify a printstream, the debug output goes to the output file you specified with `DriverManager.setLogStream()`.

Turning off debugging in your application

To turn off debugging, add this call:

```
sybdebug.debug(false);
```


Setting the CLASSPATH for debugging

Before you run your debug-enabled application, redefine the CLASSPATH environment variable to reference the */devclasses* subdirectory under your jConnect installation directory.

- For UNIX, replace `$JDBC_HOME/classes/jconn2.jar` with `$JDBC_HOME/devclasses/jconn2d.jar`.
- For Windows, replace `%JDBC_HOME%\classes\jconn2.jar` with `%JDBC_HOME%\devclasses\jconn2d.jar`.

Using the Debug methods

To customize the debugging process, you can add calls to other Debug methods.

In these methods, the first (object) parameter is usually *this* to specify the calling object. If any of these methods are static, use *null* for the object parameter.

- `println()`

Use this method to define the message to print in the output log if debugging is enabled and the object is included in the list of classes to debug. The debug output goes to the file you specified with `sybdebug.debug()`.

The syntax is:

```
sybdebug.println(object,message string);
```

For example:

```
sybdebug.println(this,"Query: "+ query);
```

produces a message similar to this in the output log:

```
myApp(thread[x,y,z]): Query: select * from authors
```

- `assert()`

Use this method to assert a condition and throw a runtime exception when the condition is not met. You can also define the message to print in the output log if the condition is not met. The syntax is:

```
sybdebug.assert(object,boolean condition,message string);
```

For example:

```
sybdebug.assert(this, amount<=buf.length, amount+"
too big!");
```

produces a message similar to this in the output log if “amount” exceeds the value of buf.length:

```
java.lang.RuntimeException:myApp(thread[x,y,z]):
Assertion failed: 513 too big!
at jdbc.sybase.utils.sybdebug.assert(
sybdebug.java:338)
at myApp.myCall(myApp.java:xxx)
at .... more stack:
```

- `startTimer()`
`stopTimer()`

Use these methods to start and stop a timer that measures the milliseconds that elapse during an event. The method keeps one timer per object, and one for all static methods. The syntax to start the timer is:

```
sybdebug.startTimer(object);
```

The syntax to stop the timer is:

```
sybdebug.stopTimer(object,message string);
```

For example:

```
sybdebug.startTimer(this);
stmt.executeQuery(query);
sybdebug.stopTimer(this,"executeQuery");
```

produces a message similar to this in the output log:

```
myApp(thread[x,y,z]):executeQuery elapsed time =
25ms
```

Capturing TDS communication

Tabular Data Stream (TDS) is Sybase’s proprietary protocol for handling communication between a client application and Adaptive Server. jConnect includes a `PROTOCOL_CAPTURE` connection property that allows you to capture raw TDS packets to a file.

If you are having problems with an application that you cannot resolve within either the application or the server, you can use `PROTOCOL_CAPTURE` to capture the communication between the client and the server in a file. You can then send the file, which contains binary data and is not directly interpretable, to Sybase Technical Support for analysis.

Note You can also use the Ribo utility to capture, translate, and display the protocol stream flowing between the client and the server. For details on how to obtain and use Ribo, visit the jConnect utilities Web page at <http://www.sybase.com/detail/1.6904.1009793.00.html>.

PROTOCOL_CAPTURE connection property

Use the `PROTOCOL_CAPTURE` connection property to specify a file for receiving the TDS packets exchanged between an application and an Adaptive Server. `PROTOCOL_CAPTURE` takes effect immediately so that TDS packets exchanged during connection establishment are written to the specified file. All packets continue to be written to the file until `Capture.pause()` is executed or the session is closed.

The following example shows the use of `PROTOCOL_CAPTURE` to send TDS data to the file *tds_data*:

```
...
props.put("PROTOCOL_CAPTURE", "tds_data");
conn = DriverManager.getConnection(url, props);
```

where *url* is the connection URL and *props* is a `Properties` object for specifying connection properties.

pause() and *resume()* methods in the *Capture* class

The `Capture` class is contained in the `com.sybase.jdbcx` package. It contains two public methods:

- `public void pause()`
- `public void resume()`

`Capture.pause()` stops the capture of raw TDS packets into a file; `Capture.resume()` restarts the capture.

The TDS capture file for an entire session can become very large. If you want to limit the size of the capture file, and you know where in an application you want to capture TDS data, you can do the following:

- 1 Immediately after you have established a connection, get the Capture object for the connection and use the `pause()` method to stop capturing TDS data:

```
Capture cap = ((SybConnection)conn).getCapture();
cap.pause();
```
- 2 Place `cap.resume()` just before the point where you want to start capturing TDS data.
- 3 Place `cap.pause()` just after the point where you want to stop capturing data.

Unsuccessful connection errors

This section addresses problems that may arise when you are trying to establish a connection or start a gateway.

Gateway connection refused

```
Gateway connection refused:
HTTP/1.0 502 Bad Gateway|Restart Connection
```

This error message indicates that something is wrong with the *hostname* or *port#* used to connect to your Adaptive Server. Check the [query] entry in *\$SYBASE/interfaces* (UNIX) or in *%SYBASE%\ini\sql.ini* (Windows).

If the problem persists after you have verified the *hostname* and *port#*, you can learn more by starting the HTTP server using the “verbose” system property.

For Windows, go to a DOS prompt and enter:

```
httpd -Dverbose=1 > filename
```

For UNIX, enter:

```
sh httpd.sh -Dverbose=1 > filename &
```

where *filename* is the debug messages output file.

Your Web server probably does not support the connect method. Applets can connect only to the host from which they were downloaded.

The HTTP gateway and your Web server must run on the same host. In this scenario, your applet can connect to the same machine/host through the port controlled by the HTTP gateway, which routes the request to the appropriate database.

To see how this is accomplished, review the source of *Isql.java* and *gateway.html* in the *sample2* subdirectory under the jConnect installation directory. Search for “proxy.”

Unable to connect to a 4.9.2 SQL Server

jConnect uses TDS 5.0 (Sybase transfer protocol). SQL Server 4.9.x uses TDS 4.6, which is not compatible with TDS 5.0.

SQL Server 10.0.2 or later is required for use with jConnect.

Memory usage in jConnect applications

The following situations and their solutions may be helpful if you notice increased memory use in jConnect applications.

- In jConnect applications, you should explicitly close all `Statement` objects and subclasses (for example, `PreparedStatement`, `CallableStatement`) after their last use to prevent statements from accumulating in memory. Closing the `ResultSet` is not sufficient.

For example:

```
ResultSet rs = _conn.prepareCall(_query).execute();  
...  
rs.close();
```

will cause problems. Instead use:

```
PreparedStatement ps = _conn.prepareCall(_query);  
ResultSet rs = ps.execute();  
...  
ps.close();  
rs.close();
```

- jConnect uses the Tabular Data Stream (TDS)—Sybase’s proprietary protocol—to communicate with Sybase database servers. As of jConnect 5.5, TDS does not support scrollable cursors. To support scrollable cursors, jConnect caches the row data on demand, on the client, on each call to `ResultSet.next()`. However, when the end of the result set is reached, the entire result set is stored in the client’s memory. Because this may cause a performance strain, Sybase recommends that you use `TYPE_SCROLL_INSENSITIVE` result sets only when the result set is reasonably small.

Stored procedure errors

This section addresses problems that may arise when you are trying to use jConnect and stored procedures.

RPC returns fewer output parameters than registered

`SQLState: JZ0SG - An RPC did not return as many output parameters as the application had registered for it.`

This error occurs if you call `CallableStatement.registerOutParam()` for more parameters than you have declared as “OUTPUT” parameters in the stored procedure. Make sure that you have declared all of the appropriate parameters as “OUTPUT.” Look at the line of code that reads:

```
create procedure yourproc (@p1 int OUTPUT, ...
```

Note If you receive this error while using Adaptive Server Anywhere (previously known as SQL Anywhere), upgrade to Adaptive Server Anywhere version 5.5.04 or later.

Fetch/state errors when output params returned

If a query does not return row data, then it should use the `CallableStatement.executeUpdate()` or `execute()` methods rather than the `executeQuery()` method.

As required by the JDBC standards, `jdbcConnect` throws a SQL exception if an `executeQuery()` has no result sets.

Stored procedure executed in unchained transaction mode

Sybase Error 7713 - Stored Procedure can only be executed in unchained transaction mode.

JDBC attempts to put the connection in `autocommit(true)` mode. The application can change the connection to chained mode using `Connection.setAutoCommit(false)` or by using a “set chained on” language command. This error occurs if the stored procedure was not created in a compatible mode.

To fix the problem, use:

```
sp_procmode procedure_name, "any mode"
```

Custom socket implementation error

You may receive an exception similar to the following while trying to set up an SSL socket when calling `sun.security.ssl.SSLSocketImpl.setEnabledCipherSuites()`:

```
java.lang.IllegalArgumentException:  
SSL_SH_anon_EXPORT_WITH_RC4_40_MDS
```

Verify that the SSL libraries are in the system library path.

This chapter describes how to fine-tune or improve performance when working with jConnect.

Topics

Improving jConnect performance

Performance tuning for prepared statements in dynamic SQL

Cursor performance

Improving jConnect performance

There are a number of ways to optimize the performance of an application using jConnect:

- Use `TextPointer.sendData()` methods to send text and image data to an Adaptive Server database. See “Updating image data in the database” on page 63.
- Create precompiled `PreparedStatement` objects for dynamic SQL statements that are used repeatedly during a session. See “Performance tuning for prepared statements in dynamic SQL” on page 114.
- Batch updates improve performance by reducing network traffic; specifically, all queries are sent to the server in one group and all responses returned to the client are sent in one group. See “Support for batch updates” on page 59.
- For sessions that are likely to move image data, large row sets, and lengthy text data, use the `PACKETSIZE` connection property to set the maximum feasible packet size.
- For TDS-tunneled HTTP, set the maximum TDS packet size and configure your Web server to support the HTTP1.1 Keep-Alive feature. Also set the `SkipDoneProc` servlet argument to true.

- Use protocol cursors, the default setting of the `LANGUAGE_CURSOR` connection property. See “`LANGUAGE_CURSOR` connection property” on page 121 for more information.
- If you use `TYPE_SCROLL_INSENSITIVE` result sets, only use them when the result set is reasonably small. See “Support for `SCROLL_INSENSITIVE` result sets in jConnect” on page 56 for more information.

Additional considerations for improving performance are described below.

BigDecimal rescaling

The JDBC 1.0 specification requires a scale factor with `getBigDecimal()`. Then, when a `BigDecimal` object is returned from the server, it must be rescaled using the original scale factor you used with `getBigDecimal()`.

To eliminate the processing time required for rescaling, use the JDBC 2.0 `getBigDecimal()` method, which jConnect implements in the `SybResultSet` class and does not require a *scale* value:

```
public BigDecimal getBigDecimal(int columnIndex)
    throws SQLException
```

For example:

```
SybResultSet rs =
    (SybResultSet)stmt.executeQuery("SELECT
    numeric_column from T1");
while (rs.next())
{
    BigDecimal bd rs.getBigDecimal(
        "numeric_column");
    ...
}
```

REPEAT_READ connection property

You can improve performance on retrieving a result set from the database if you set the `REPEAT_READ` connection property to false. However, when `REPEAT_READ` is “false:”

- You must read column values in order, according to column index. This is difficult if you want to access columns by name rather than column number.
- You cannot read a column value in a row more than once.

Character-set conversion

SunloConverter

If you are using multibyte character sets and need to improve driver performance, you can use the `SunloConverter` class provided with the `jConnect` samples. This converter is based on the `sun.io` classes provided by the Java Software Division of Sun Microsystems, Inc.

The `SunloConverter` class is not a pure Java implementation of the character-set converter feature, and therefore is not integrated with the standard `jConnect` product. However, we have provided this converter class for your reference, and you can use it with the `jConnect` driver to improve character-set conversion performance.

Note Based on Sybase testing, the `SunloConverter` class improved performance on all VMs on which it was tested. However, the Java Software Division of Sun Microsystems, Inc. reserves the right to remove or change the `sun.io` classes with future releases of the JDK, and therefore this `SunloConverter` class may not be compatible with later JDK releases.

To use the `SunloConverter` class, you must install the `jConnect` sample applications. See the *Sybase jConnect for JDBC Installation Guide* for complete instructions on installing `jConnect` and its components, including the sample applications. Once the samples are installed, set the `CHARSET_CONVERTER_CLASS` connection property to reference the `SunloConverter` class in the *sample2* subdirectory under your `jConnect` installation directory.

If you are using a database with its default character set as `iso_1`, or if you are using only the first 7 bits of ASCII, you can gain significant performance benefits by using the `TruncationConverter`. See “`jConnect` character-set converters” on page 34.

Performance tuning for prepared statements in dynamic SQL

In Embedded SQL, dynamic statements are SQL statements that need to be compiled at runtime, rather than statically. Typically, dynamic statements contain input parameters, although this is not a requirement. In SQL, the prepare command is used to precompile a dynamic statement and save it so that it can be executed repeatedly without being recompiled during a session.

If a statement is used multiple times in a session, precompiling it provides better performance than sending it to the database and compiling it for each use. The more complex the statement, the greater the performance benefit.

If a statement is likely to be used only a few times, precompiling it may be inefficient because of the overhead involved in precompiling, saving, and later deallocating it in the database.

Precompiling a dynamic SQL statement for execution and saving it in memory uses time and resources. If a statement is not likely to be used multiple times during a session, the costs of doing a database prepare may outweigh its benefits. Another consideration is that once a dynamic SQL statement is prepared in the database, it is very similar to a stored procedure. In some cases, it may be preferable to create stored procedures and have them reside on the server, rather than defining prepared statements in the application. This is discussed under “Choosing prepared statements and stored procedures” on page 115.

You can use `jConnect` to optimize the performance of dynamic SQL statements on a Sybase database as follows:

- Create `PreparedStatement` objects that contain precompiled statements in cases where a statement is likely to be executed several times in a session.
- Create `PreparedStatement` objects that contain uncompiled SQL statements in cases where a statement will be used very few times in a session.

As described in the following sections, the optimal way to set the `DYNAMIC_PREPARE` connection property and create `PreparedStatement` objects is likely to depend on whether your application needs to be portable across JDBC drivers or whether you are writing an application that allows `jConnect`-specific extensions to JDBC.

`jConnect` 4.1 and later provide performance tuning features for dynamic SQL statements.

Choosing prepared statements and stored procedures

If you create a `PreparedStatement` object containing a precompiled dynamic SQL statement, once the statement is compiled in the database, it effectively becomes a stored procedure that is retained in memory and attached to the data structure associated with your session. In deciding whether to maintain stored procedures in the database or to create `PreparedStatement` objects containing compiled SQL statements in your application, resource demands and database and application maintenance are important considerations:

- Once a stored procedure is compiled, it is globally available across all connections. In contrast, a dynamic SQL statement in a `PreparedStatement` object needs to be compiled and deallocated in every session that uses it.
- If your application accesses multiple databases, using stored procedures means that the same stored procedures need to be available on all target databases. This can create a database maintenance problem. If you use `PreparedStatement` objects for dynamic SQL statements, you avoid this problem.
- If your application creates `CallableStatement` objects for invoking stored procedures, you can encapsulate SQL code and table references in the stored procedures. You can then modify the underlying database or SQL code without have to change the application.

Prepared statements in portable applications

If your application is to run on databases from different vendors and you want some `PreparedStatement` objects to contain precompiled statements and others to contain uncompiled statements, proceed as follows:

- When you access a Sybase database, make sure that the `DYNAMIC_PREPARE` connection property is set to true.
- To return `PreparedStatement` objects containing precompiled statements, use `Connection.prepareStatement()` in the standard way:

```
PreparedStatement ps_precomp =  
    Connection.prepareStatement(sql_string);
```

- To return `PreparedStatement` objects containing uncompiled statements, use `Connection.prepareCall()`.

`Connection.prepareCall()` returns a `CallableStatement` object, but since `CallableStatement` is a subclass of `PreparedStatement`, you can upcast a `CallableStatement` object to a `PreparedStatement` object, as in the following example:

```
PreparedStatement ps_uncomp =  
    Connection.prepareCall(sql_string);
```

The `PreparedStatement` object *ps_uncomp* is guaranteed to contain an uncompiled statement, since only `Connection.prepareStatement()` is implemented to return `PreparedStatement` objects containing precompiled statements.

Prepared statements with jConnect extensions

If you are not concerned about portability across drivers, you can write code that uses `SybConnection.prepareStatement()` to specify whether a `PreparedStatement` object contains precompiled or uncompiled statements. In this case, how you code prepared statements is likely to depend on whether most of the dynamic statements in an application are likely to be executed many times or only a few times during a session.

If most dynamic statements are executed very few times

For an application in which most dynamic SQL statements are likely to be executed only once or twice in a session:

- Set the connection property `DYNAMIC_PREPARE` to false.
- To return `PreparedStatement` objects containing uncompiled statements, use `Connection.prepareStatement()` in the standard way:

```
PreparedStatement ps_uncomp =  
    Connection.prepareStatement(sql_string);
```

- To return `PreparedStatement` objects containing precompiled statements, use `SybConnection.prepareStatement()` with *dynamic* set to “true:”

```
PreparedStatement ps_precomp =  
    (SybConnection)conn.prepareStatement(sql_string, true);
```

If most dynamic statements are executed many times in a session

If most of the dynamic statements in an application are likely to be executed many times in the course of a session, proceed as follows:

- Set the connection property `DYNAMIC_PREPARE` to true.
- To return `PreparedStatement` objects containing precompiled statements, use `Connection.prepareStatement()` in the standard way:

```
PreparedStatement ps_precomp =
    Connection.prepareStatement( sql_string );
```

- To return `PreparedStatement` objects containing uncompiled statements, you can use either `Connection.prepareCall()` (see the third bullet under Prepared statements in portable applications) or `SybConnection.prepareStatement()`, with *dynamic* set to “false:”

```
PreparedStatement ps_uncomp =
    (SybConnection)conn.prepareStatement( sql_string,
    false );
```

```
PreparedStatement ps_uncomp =
    Connection.prepareCall( sql_string );
```

Connection.prepareStatement()

`jConnect` implements `Connection.prepareStatement()` so you can set it to return either precompiled SQL statements or uncompiled SQL statements in `PreparedStatement` objects. If you set `Connection.prepareStatement()` to return precompiled SQL statements in `PreparedStatement` objects, it sends dynamic SQL statements to the database to be precompiled and saved exactly as they would be under direct execution of the prepare command. If you set `Connection.prepareStatement()` to return uncompiled SQL statements, it returns them in `PreparedStatement` objects without sending them to the database.

The type of SQL statement that `Connection.prepareStatement()` returns is determined by the connection property `DYNAMIC_PREPARE`, and applies throughout a session.

For Sybase-specific applications, `jConnect 5.5` provides a `prepareStatement()` method under the `jConnect SybConnection` class.

`SybConnection.prepareStatement()` allows you to specify whether an individual dynamic SQL statement is to be precompiled, independent of the session-level setting of the `DYNAMIC_PREPARE` connection property.

DYNAMIC_PREPARE connection property

DYNAMIC_PREPARE is a Boolean-valued connection property for enabling dynamic SQL prepared statements:

- If DYNAMIC_PREPARE is set to true, every invocation of `Connection.prepareStatement()` during a session attempts to return a precompiled statement in a `PreparedStatement` object.

In this case, when a `PreparedStatement` is executed, the statement it contains is already precompiled in the database, with place holders for dynamically assigned values, and the statement needs only to be executed.

- If DYNAMIC_PREPARE is set to false for a connection, the `PreparedStatement` object returned by `Connection.prepareStatement()` does not contain a precompiled statement.

In this case, each time a `PreparedStatement` is executed, the dynamic SQL statement it contains must be sent to the database to be both compiled and executed.

The default value for DYNAMIC_PREPARE is false.

In the following example, DYNAMIC_PREPARE is set to true to enable precompilation of dynamic SQL statements. In the example, `props` is a `Properties` object for specifying connection properties.

```
...
props.put("DYNAMIC_PREPARE", "true")
Connection conn = DriverManager.getConnection(url,
props);
```

When DYNAMIC_PREPARE is set to true, note that:

- Not all dynamic statements can be precompiled under the prepare command. The SQL-92 standard places some restrictions on the statements that can be used with the prepare command, and individual database vendors may have their own constraints.
- If the database generates an error because it is unable to precompile and save a statement sent to it through `Connection.prepareStatement()`, `jdbcConnect` traps the error and returns a `PreparedStatement` object containing an uncompiled dynamic SQL statement. Each time the `PreparedStatement` object is executed, the statement is re-sent to the database to be compiled and executed.

- A precompiled statement resides in memory in the database and persists either to the end of a session or until its `PreparedStatement` object is explicitly closed. Garbage collection on a `PreparedStatement` object does not remove the prepared statement from the database.

As a general rule, you should explicitly close every `PreparedStatement` object after its last use to prevent prepared statements from accumulating in server memory during a session and slowing performance.

SybConnection.prepareStatement()

If your application allows `jConnect`-specific extensions to JDBC, you can use the `SybConnection.prepareStatement()` extension method to return dynamic SQL statements in `PreparedStatement` objects:

```
PreparedStatement SybConnection.prepareStatement  
(String sql_stmt,  
    boolean dynamic) throws SQLException
```

`SybConnection.prepareStatement()` can return `PreparedStatement` objects containing either precompiled or uncompiled SQL statements, depending on the setting of the *dynamic* parameter. If *dynamic* is true, `SybConnection.prepareStatement()` returns a `PreparedStatement` object with a precompiled SQL statement. If *dynamic* is false, it returns a `PreparedStatement` object with an uncompiled SQL statement.

The following example shows the use of

`SybConnection.prepareStatement()` to return a `PreparedStatement` object containing a precompiled statement:

```
PreparedStatement precomp_stmt =  
    ((SybConnection) conn).prepareStatement( "SELECT *  
FROM  
    authors WHERE au_fname LIKE ?", true);
```

In the example, the connection object *conn* is downcast to a `SybConnection` object to allow the use of `SybConnection.prepareStatement()`. The SQL string passed to `SybConnection.prepareStatement()` will be precompiled in the database, even if the connection property `DYNAMIC_PREPARE` is false.

If the database generates an error because it is unable to precompile a statement sent to it through `SybConnection.prepareStatement()`, `jConnect` throws a `SQLException` and the call fails to return a `PreparedStatement` object. This is unlike `Connection.prepareStatement()`, which traps SQL errors and, in the event of an error, returns a `PreparedStatement` object containing an uncompiled statement.

ESCAPE_PROCESSING_DEFAULT connection property

By default `jConnect` parses all SQL statements submitted to the database for valid JDBC function escapes. If your application is not going to use JDBC function escapes in its SQL calls, you can set this connection property to *false* to circumvent this parsing. This may give a slight performance benefit.

Cursor performance

When you use the `Statement.setCursorName()` method or the `setFetchSize()` method in the `SybCursorResultSet` class, `jConnect` creates a cursor in the database. Other methods cause `jConnect` to open, fetch, and update a cursor.

Versions of `jConnect` earlier than 4.0 can create and manipulate cursors only by sending SQL statements with explicit cursor commands to the database for parsing and compilation.

`jConnect` version 4.0 and later create and manipulate cursors either by sending SQL statements to the database or by encoding cursor commands as tokens within the Tabular Data Stream (TDS) communication protocol. Cursors of the first type are “language cursors;” cursors of the second type are “protocol cursors.”

Protocol cursors provide better performance than language cursors. In addition, not all databases support language cursors. For example, Adaptive Server Anywhere databases do not support language cursors.

In `jConnect`, the default condition is for all cursors to be protocol cursors. However, the `LANGUAGE_CURSOR` connection property gives you the option of having cursors created and manipulated through language commands in the database.

LANGUAGE_CURSOR connection property

LANGUAGE_CURSOR is a Boolean-valued connection property in jConnect that allows you to determine whether cursors are created as protocol cursors or language cursors:

- If LANGUAGE_CURSOR is set to false, all cursors created during a session are protocol cursors, which provide better performance. jConnect creates and manipulates the cursors by sending cursor commands as tokens in the TDS protocol.

By default, LANGUAGE_CURSOR is set to false.

- If LANGUAGE_CURSOR is set to true, all cursors created during a session are language cursors. jConnect creates and manipulates the cursors by sending SQL statements to the database for parsing and compilation.

There is no known advantage to setting LANGUAGE_CURSOR to true, but the option is provided in case an application displays unexpected behavior when LANGUAGE_CURSOR is false.

This chapter explains how to migrate applications that use Sybase extensions from jConnect version 4.0 or earlier to use jConnect versions 4.1 and later.

Topics

Migrating applications to jConnect 4.5 and 5.5

Sybase extensions

Migrating jConnect applications

Migrating applications to jConnect 4.5 and 5.5

If you upgrade to jConnect 4.5 or 5.5 from earlier versions, the following table shows which upgrade paths require you to make changes and recompile the source code.

Legend:

- A Change CLASSPATH for new installation structure
- B Recompile to use new jConnect 5.x driver
- C Verify that the new driver is first in your CLASSPATH.

See below for additional details.

Upgrading from	To jConnect version		
jConnect version	4.5	5.2	5.5
4.2/4.5	A	BC	BC
5.2	-	-	A

A. Use the new Sybase extensions.

- 1 Change package imports from
import com.sybase.jdbc.*

A. Change CLASSPATH for new JDBC_HOME installation structure

to

```
import com.sybase.jdbcx.*;
```

2 Use new Sybase extensions APIs. See “Sybase extensions” on page 125.

Set JDBC_HOME to the top directory of the jConnect driver you installed.
For example:

For jConnect 4.5:

```
JDBC_HOME=jConnect-4_5
```

For jConnect 5.5:

```
JDBC_HOME=<jConnect installation directory>
```

For more information on setting JDBC_HOME, see “Setting Environment Variables” in Chapter 1 of the *jConnect for JDBC Installation Guide*.

Version change		CLASSPATH includes
From	4.2/4.5	<i>JDBC_HOME/classes</i>
To	5.2	<i>JDBC_HOME/jconn2.jar</i>
From	4.2	<i>JDBC_HOME/classes</i>
To	4.5	<i>JDBC_HOME/classes</i>
From	5.2	<i>JDBC_HOME/classes/jconn2.jar</i>
To	5.5	<i>JDBC_HOME/classes/jconn2.jar</i>

B. Recompile to use new jConnect 5.x driver

Change the source code where the driver is loaded from:

```
Class.forName( "com.sybase.jdbc.SybDriver" );
```

to

```
Class.forName( "com.sybase.jdbc2.jdbc.SybDriver" );
```

C. Verify that new jConnect driver is in CLASSPATH

To verify that the new jConnect driver is first in your CLASSPATH, display your CLASSPATH and find the directory where you installed the new jConnect driver.

Verify that the classes for the new driver (in the *classes/jconn2.jar* for 5.x) appear in the CLASSPATH before the classes for any previously installed driver.

Sybase extensions

jConnect version 4.1 and later include the package `com.sybase.jdbcx` that contains all of the Sybase extensions to JDBC. In versions of jConnect previous to 4.1, these extensions were available in the `com.sybase.jdbc` and `com.sybase.utils` packages.

`com.sybase.jdbcx` provides a consistent interface across different versions of jConnect. All of the Sybase extensions are defined as Java interfaces, which allows the underlying implementations to change without affecting applications built using these interfaces.

When you develop new applications that use Sybase extensions, use `com.sybase.jdbcx`. The interfaces in this package allow you to upgrade applications to versions of jConnect that follow version 4.0 with minimal changes.

Note Applications previously built using the Sybase extensions to the JDBC API, which were available in `com.sybase.jdbc` and `com.sybase.utils`, will continue to work under jConnect 4.x; however, all Sybase extensions in `com.sybase.jdbc` and `com.sybase.utils` have been marked deprecated.

Some of the Sybase extensions have been changed to accommodate the new `com.sybase.jdbcx` interface.

Change example

If an application uses the `SybMessageHandler`, the code differences would be:

- **jConnect 4.0** code:

```
import com.sybase.jdbc.SybConnection;
import com.sybase.jdbc.SybMessageHandler;
.
.
Connection con = DriverManager.getConnection(url, props);
SybConnection sybCon = (SybConnection) con;
sybCon.setMessageHandler(new ConnectionMsgHandler());
```

- **jConnect 4.1** and later code:

```
import com.sybase.jdbcx.SybConnection;
import com.sybase.jdbcx.SybMessageHandler;
.
```

```

Connection con = DriverManager.getConnection(url, props);
SybConnection sybCon = (SybConnection) con;
sybCon.setSybMessageHandler(new ConnectionMsgHandler());

```

See the samples provided with jConnect for more examples of how to use Sybase extensions.

Method names

The following table lists how methods were been renamed in the new interface.

Class	Old name	New name
SybConnection	getCapture()	createCapture()
SybConnection	setMessageHandler()	setSybMessageHandler()
SybConnection	getMessageHandler()	getSybMessageHandler()
SybStatement	setMessageHandler()	setSybMessageHandler()
SybStatement	getMessageHandler()	getSybMessageHandler()

Debug class

Direct static references to the Debug class are no longer supported, but exist in deprecated form in the com.sybase.utils package. To use jConnect debugging facilities, use the getDebug() method of the SybDriver class to obtain a reference to the Debug class. For example:

```

import com.sybase.jdbcx.SybDriver;
import com.sybase.jdbcx.Debug;
.
.
.
SybDriver sybDriver =
    (SybDriver)Class.forName
        ("com.sybase.jdbc2.jdbc.SybDriver") newInstance();
Debug sybDebug = sybDriver.getDebug();
sybDebug.debug(true, "ALL", System.out);

```

A complete list of Sybase extensions is in the jConnect javadoc documentation located in the *docs/* directory of your jConnect installation directory.

This chapter describes Web server gateways and explains how to use them with jConnect.

TOPICS

About Web server gateways

Using the TDS-tunnelling servlet

About Web server gateways

If your database server runs on a different host than your Web server, or if you are developing Internet applications that must connect to a secure database server through a firewall, you need a gateway to act as a proxy, providing a path to the database server.

To connect to servers using the Secure Sockets Layer (SSL) protocol, jConnect provides a Java servlet that you can install on any Web server that supports the `javax.servlet` interfaces. This servlet enables jConnect to support encryption using the Web server as the gateway.

Note jConnect includes support for SSL on the client system. For information on jConnect's client-side support of SSL, see "Implementing custom socket plug-ins" on page 28.

TDS tunnelling

jConnect uses TDS to communicate with database servers. HTTP-tunnelled TDS is useful for forwarding requests. Requests from a client to a back-end server that go through the gateway contain TDS in the body of the request. The request header indicates the length of the TDS included in the request packet.

TDS is a connection-oriented protocol, whereas HTTP is not. To support security features such as encryption for Internet applications, jConnect uses a TDS-tunnelling servlet to maintain a logical connection across HTTP requests. The servlet generates a session ID during the initial login request, and the session ID is included in the header of every subsequent request. Using session IDs lets you identify active sessions, and even resume a session as long as the servlet has an open connection using that specific session ID.

The logical connection provided by the TDS-tunnelling servlet enables jConnect to support encrypted communication between two systems—for example, a jConnect client with the `CONNECT_PROTOCOL` connection property set to “https” connecting to a Web server running the TDS-tunnelling servlet.


jConnect and gateway configuration

There are several options for setting up your Web servers and Adaptive Servers. Four common configurations are described below. These examples show where to install the jConnect driver and when to use a gateway with the TDS-tunnelling servlet.

Web server and Adaptive Server on one host

In this two-tier configuration, the Web server and Adaptive Server are both installed on the same host.

- Install jConnect on the Web server host.
- No gateway required.

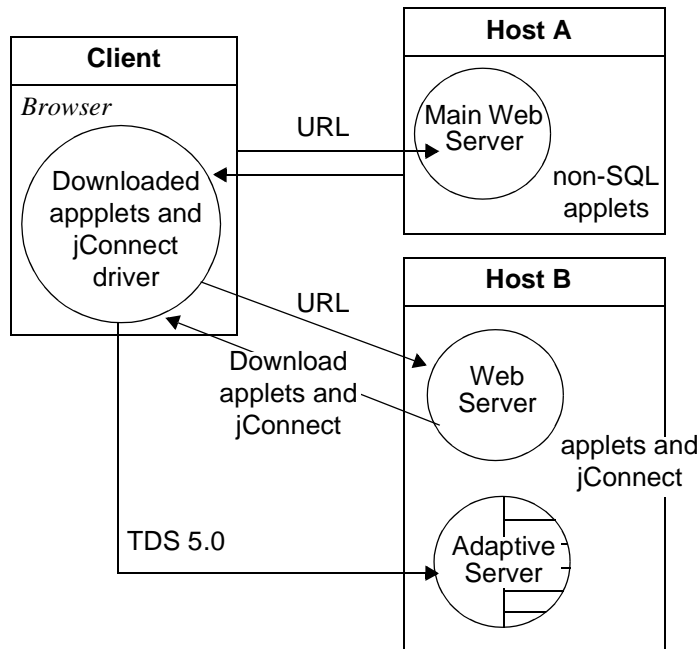
Figure 6-1: Web server and Adaptive Server on one host

Dedicated JDBC Web server and Adaptive Server on one host

With this configuration, you have a separate host for your main Web server. A second host is shared by a Web server specifically for Adaptive Server access and the Adaptive Server. Links from the main server direct requests requiring SQL access to the dedicated Web server.

- Install jConnect on the second (Adaptive Server) host.
- No gateway required.

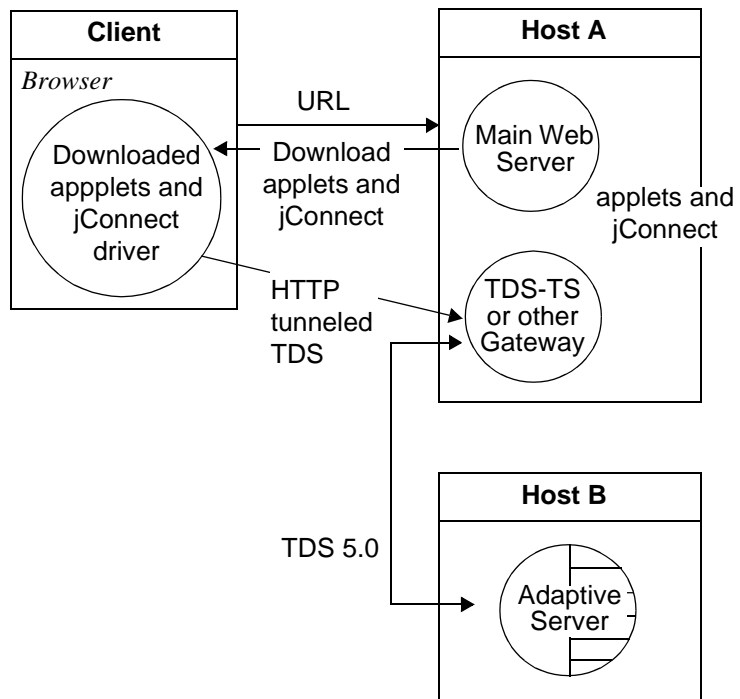
Figure 6-2: Main Web server on separate host



Web server and Adaptive Server on separate hosts

In this three-tier configuration, the Adaptive Server is on a separate host from the Web server. jConnect requires a gateway to act as a proxy to the Adaptive Server.


- Install jConnect on the Web server host.
- Requires a TDS-tunnelling servlet or a different gateway.

Figure 6-3: Web Server and Adaptive Server on separate hosts

Connecting to a server through a firewall

To connect to a server protected by a firewall, you must use a Web server with the TDS-tunnelling servlet to support transmission of database request responses over the Internet.

- Install jConnect on the Web server host.
- Requires a Web server that supports the `javax.servlet` interfaces.

Figure 6-4: Connect a server through a firewall

Usage requirements

Reading the *index.html* file

Use your Web browser to view the *index.html* file in your jConnect installation directory. *index.html* provides links to the jConnect documentation and sample code.

Note If you use Netscape on the same machine where you have installed jConnect, make sure that your browser does not have access to your CLASSPATH environment variable. See “Restrictions on Setting CLASSPATH When You Use Netscape” in Chapter 3 of the *Sybase jConnect for JDBC Installation Guide and Release Bulletin*.

- 1 Open your Web browser.
- 2 Enter the URL that matches your setup. For example, if your browser and the Web server are running on the same host, enter:

`http://localhost:8000/index.html`

If the browser and the Web server are running on different hosts, enter:

`http://host:port/index.html`

where *host* is the name of the host on which the Web server is running, and *port* is the listening port.

Running the sample Isql applet

After loading the *index.html* file in your browser:

- 1 Click “Run Sample JDBC Applets.”
This takes you to the jConnect Sample Programs page.
- 2 Move down the Sample Programs page to find the table under “Executable Samples.”
- 3 Locate “Isql.java” in the table and click Run at the end of the row.

The sample Isql.java applet prompts for a simple query on a sample database and displays the results. The applet displays a default Adaptive Server host name, port number, user name (*guest*), password (*sybase*), database, and query. Using the default values, the applet connects to the Sybase demonstration database. It returns results after you click Go.

Troubleshooting

Under UNIX, if the applet does not appear as expected, you can modify the applet screen dimensions:

- 1 Use a text editor to edit the following:

\$JDBC_HOME/sample2/gateway.html

- 2 Change the height parameter in line 7 to 650. You can experiment with different height settings.
- 3 Reload the Web page on your browser.

Using the TDS-tunnelling servlet

To use the TDS-tunnelling servlet, you need a Web server that supports the javax.servlet interfaces, such as Sun Microsystems, Inc.'s Java Web server. When you install the Web server, include the jConnect TDS-tunnelling servlet in the list of active servlets. You can also set servlet parameters to define connection timeouts and maximum packet size.

With the TDS-tunnelling servlet, requests from a client to the back-end server that go through the gateway include a GET or POST command, the TDS session ID (after the initial request), back-end address, and status of the request.

TDS is in the body of the request. Two header fields indicate the length of the TDS stream and the session ID assigned by the gateway.

When the client sends a request, the Content-Length header field indicates the size of the TDS content, and the request command is POST. If there is no TDS data in the request because the client is either retrieving the next portion of the response data from the server, or closing the connection, the request command is GET.

The following example illustrates how information is passed between the client and an HTTPS gateway using the TDS-tunneled HTTPS protocol; it shows a connection to a back-end server named DBSERVER with a port number "1234."

Table 6-1: Client to gateway login request. No session ID.

<i>Query</i>	POST/tds?ServerHost=dbserver&ServerPort=1234&Operation=more HTTP/1.0
<i>Header</i>	Content-Length: 605
<i>Content</i> (TDS)	Login request

Table 6-2: Gateway to client. Header contains session ID assigned by the TDS servlet.

<i>Query</i>	200 SUCCESS HTTP/1.0
<i>Header</i>	Content-Length: 210 TDS-Session: TDS00245817298274292
<i>Content</i> (TDS)	Login acknowledgment EED

Table 6-3: Client to gateway. Headers for all subsequent requests contain the session ID.

<i>Query</i>	POST/tds?TDS-Session=TDS00245817298274292&Operation=more HTTP/1.0
<i>Header</i>	Content-Length: 32
<i>Content</i> (TDS)	Query "SELECT * from authors"

Table 6-4: Gateway to client. Headers for all subsequent responses contain the session ID.

<i>Query</i>	200 SUCCESS HTTP/1.0
<i>Header</i>	Content-Length: 2048 TDS-Session: TDS00245817298274292
<i>Content</i> (TDS)	Row format and some rows from query response

TDS-tunnelling servlet system requirements

To use the jConnect servlet for TDS-tunneled HTTP, you need:

- A Web server that supports javax.servlet interfaces. To install the server, follow the instructions that are provided with it.
- A Web browser that supports JDK 1.1, such as Netscape 4.0, Internet Explorer 4.0, or HotJava.

Installing the servlet

Your jConnect installation includes a *gateway2* subdirectory under the *classes* directory. The subdirectory contains files required for the TDS-tunnelling servlet.

Copy the jConnect gateway package to a *gateway2* subdirectory under your Web server's servlets directory. Once you have copied the servlets, activate the servlets by following the instructions for your Web server.

Setting servlet arguments

When you add the servlet to your Web server, you can enter optional arguments to customize performance:

- *SkipDoneProc* [*true/false*] – Sybase databases often return row count information as intermediate processing steps are performed during the execution of a query. Usually client applications ignore this data. If you set *SkipDoneProc* to true, the servlet will remove this extra information from responses “on the fly,” which reduces network usage and processing requirements on the client. This is particularly useful when using HTTPS/SSL because the unwanted data does not get encrypted/decrypted before it is ignored.
- *TdsResponseSize* – set the maximum TDS packet size for the tunneled HTTPS. A larger *TdsResponseSize* is more efficient if you have only a few users with a large volume of data. Use a smaller *TdsResponseSize* if you have many users making smaller transactions.
- *TdsSessionIdleTimeout* – define the amount of time (in milliseconds) that the server connection can remain idle before the connection is automatically closed. The default *TdsSessionIdleTimeout* is 600,000 (10 minutes).

If you have interactive client programs that may be idle for long periods of time and you do not want the connections broken, increase the *TdsSessionIdleTimeout*.

You can also set the connection timeout value from the jConnect client using the `SESSION_TIMEOUT` connection property. This is useful if you have specific applications that may be idle for longer periods. In this case, set a longer timeout for those connections with the `SESSION_TIMEOUT` connection property, rather than setting it for the servlet.

- *Debug* – turn on debugging. See “Debugging with jConnect” on page 101.

Enter the servlet arguments in a comma-delimited string. For example:

```
TdsResponseSize=[size],TdsSessionIdleTimeout=  
[timeout],Debug=true
```

Refer to your Web server documentation for complete instructions on entering servlet arguments.

Invoking the servlet

jConnect determines when to use the gateway where the TDS-tunnelling servlet is installed based on the path extension of the *proxy* connection property. jConnect recognizes the servlet path extension to the *proxy* and invokes the servlet on the designated gateway.

Define the connection URL using this format:

```
http://host:port/TDS-servlet-path
```

jConnect invokes the TDS-tunnelling servlet on the Web server to tunnel TDS through HTTP. The servlet path must be the path you defined in your Web server's servlet alias list.

Tracking active TDS sessions

You can view information about active TDS sessions, including the server connections for each session. Use your Web browser to open the administrative URL:

```
http://host:port/TDS-servlet-path?Operation=list
```

For example, if your server is MYSERVER and the TDS servlet path is */tds*, enter:

```
http://myserver:8080/tds?Operation=list
```

This shows you a list of active TDS sessions. You can click on a session to see more information, including the server connection.

Terminating TDS sessions

You can use the URL described above to terminate any active TDS session. Click on an active session from the list of sessions on the first page, then click *Terminate This Session*.

Resuming a TDS session

You can set the `SESSION_ID` connection property so that, if necessary, you can resume an existing open connection. When you specify a `SESSION_ID`, `jConnect` skips the login phase of the protocol and resumes the connection with the gateway using the designated session ID. If the session ID you specified does not exist on the servlet, `jConnect` throws a SQL exception the first time you attempt to use the connection.

TDS and Netscape Enterprise Server 3.5.1 on Solaris

Netscape Enterprise Server 3.5.1 does not support the `javax.servlet.ServletConfig.getInitParameters()` or `javax.servlet.ServletConfig.getInitParameterNames()` methods. To provide the necessary parameter values, you need to replace calls to `getInitParameter()` and `getInitParameterNames()` with hard-coded parameter values in *TDSTunnelServlet.java*.

To enter the required parameter values in *TDSTunnelServlet.java* and use TDS tunnelling with Netscape Enterprise Server 3.5.1 on Solaris:

- 1 Hard code parameter values in *TDSTunnelServlet.java*.
- 2 Create *.class* files from the class declarations in *TDSTunnelServlet.java*.

This should result in the following files:

- *TDSTunnelServlet.class*
- *TdsSession.class*
- *TdsSessionManager.class*

- 3 Create a directory for the *.class* files under your Netscape Enterprise Server 3.5.1 (NSE_3.5.1) installation directory, as follows:

```
mkdir NSE_3.5.1_install_dir/plugins/java/servlets/gateway
```

- 4 Copy the *.class* files derived from *TDSTunnelServlet.java* to the directory you just created.
- 5 Copy the classes under `$JDBC_HOME/classes/com/sybase` to `NSE_3.5.1_install_dir/docs/com/sybase`.

An easy way to do this is to recursively copy everything under `$JDBC_HOME/classes` to `NSE_3.5.1_install_dir/docs`, as:

```
cp -r $JDBC_HOME/classes NSE_3.5.1_install_dir/docs
```

This copies a number of files and directories that are not under *\$JDBC_HOME/classes/com/sybase*. The extra files and directories are harmless, but take up disk space. You can delete them to reclaim the disk space.

6 Set the *proxy* URL to the TDS-tunnelling servlet.

For example, in *\$JDBC_HOME/sample2/gateway.html*, you would edit the *proxy* parameter to appear as follows:

```
<param name=proxy value="http://hostname/servlet/  
gateway_name.TDSTunnel_Servlet_name">
```


SQL Exception and Warning Messages

The following table lists the SQL exception and warning messages that you may encounter when using jConnect.

SQL state	Message/description/action
010AF	<p>SEVERE WARNING: An assertion has failed, please use devclasses to determine the source of this serious bug. Message = ____.</p> <p>Description: An internal assertion in the jConnect driver has failed.</p> <p>Action: Use the devclasses debug classes to determine the reason for this message and report the problem to Sybase Technical Support.</p>
010DF	<p>Attempt to set database at login failed. Error message:____.</p> <p>Description: jConnect cannot connect to the database specified in the connection URL.</p> <p>Action: Be sure the database name in the URL is correct. Also, if connecting to Adaptive Server Anywhere, use the SERVICENAME connection property to specify the database.</p>
010DP	<p>Duplicate connection property ____ ignored.</p> <p>Description: A connection property is defined twice. It may be defined twice in the driver connection properties list with different capitalization, for example “password” and “PASSWORD.” Connection property names are not case-sensitive, and therefore jConnect does not distinguish between property names with the same name but different capitalization.</p> <p>The connection property may also be defined both in the connection properties list, and in the URL. In this case, the property value in the connection property list takes precedence.</p> <p>Action: Make sure your application defines the connection property only once. However, you may want you application to take advantage of the precedence of connection properties defined in the property list over those defined in the URL. In this case, you can safely ignore this warning.</p>
010HA	<p>The server denied your request to use the high-availability feature. Please reconfigure your database, or do not request a high-availability session.</p> <p>Description: The server to which jConnect attempted an HA-enabled connection did not allow the connection.</p> <p>Action: Reconfigure the server to support high availability failover or do not set REQUEST_HA_SESSION to true.</p>

SQL state	Message/description/action
010HD	<p>Sybase high-availability failover is not supported by this type of database server.</p> <p>Description: The database to which jConnect attempted a connection does not support high availability failover.</p> <p>Action: Connect only to database servers that support high availability failover.</p>
010HT	<p>Hostname property truncated, maximum length is 30.</p> <p>Description: You provided a String greater than 30 characters for the HOSTNAME connection property, or the host machine on which the jConnect application is running has a name longer than 30 bytes in length.</p> <p>Action: No action is necessary, since this is just a warning to indicate that jConnect is truncating the name to 30 bytes. However, if you wish to avoid this warning, you should set the HOSTNAME to a String less than or equal to 30 bytes in length.</p>
010MX	<p>Metadata accessor information was not found on this database. Please install the required tables as mentioned in the jConnect documentation. Error encountered while attempting to retrieve metadata information: _____</p> <p>Description: The server may not have the necessary stored procedures for returning metadata information.</p> <p>Action: Make sure that stored procedures for providing metadata are installed on the server. See “Installing Stored Procedures” in Chapter 3 of the <i>jConnect for JDBC Installation Guide</i>.</p>
010P4	<p>An output parameter was received and ignored.</p> <p>Description: The query you executed returned an output parameter but the application result-processing code did not fetch it so it was ignored.</p> <p>Action: If your application needs the output parameter data, you must rewrite the application so it can get it. This may require using a CallableStatement to execute the query, and adding calls to registerOutputParameter() and getXXX(). You can also prevent jConnect from returning this warning, and possibly get a performance improvement, by setting the DISABLE_UNPROCESSED_PARAM_WARNINGS connection property to true.</p>
010PF	<p>One or more jars specified in the PRELOAD_JARS connection property could not be loaded.</p> <p>Description: This happens when using the DynamicClassLoader with the PRELOAD_JARS connection property set to a comma-delimited list of JAR names. When the DynamicClassLoader opens its connection to the server from which the classes are to be loaded, it attempts to “preload” all the JAR files mentioned in this connection property. If one or more of the JAR names specified does not exist on the server, the above error message results.</p> <p>Action: Verify that every JAR file mentioned in your application’s PRELOAD_JARS connection property exists and is accessible on the server.</p>

SQL state	Message/description/action
010PO	<p>Property LITERAL_PARAM has been reset to "false" because DYNAMIC_PREPARE was set to "true".</p> <p>Description: If you wish to use precompiled dynamic statements, then you must allow for parameters to be sent to those statements (if the statements take parameters). Setting LITERAL_PARAMS to true forces all parameters to be sent as literal values in the SQL that you send to the server. Therefore, you cannot set both properties to true.</p> <p>Action: To avoid this warning, do not set LITERAL_PARAMS to true when you wish to use dynamic SQL. See “Performance tuning for prepared statements in dynamic SQL” on page 114 for more information.</p>
010RC	<p>The requested ResultSet type and concurrency is not supported. They have been converted.</p> <p>Description: You requested a type and concurrency combination for the ResultSet that is not supported. The requested values had to be converted. See “Using cursors with result sets” on page 48 for more information about what ResultSet types and concurrencies are available in <i>jdbc</i>.</p> <p>Action: Request a type and concurrency combination for the ResultSet that is supported.</p>
010SJ	<p>Metadata accessor information was not found on this database. Please install the required tables as mentioned in the <i>jdbc</i> documentation.</p> <p>Description: The metadata information is not configured on the server.</p> <p>Action: If your application requires metadata, install the stored procedures for returning metadata that come with <i>jdbc</i> (see “Installing Stored Procedures” in Chapter 3 of the <i>jdbc for JDBC Installation Guide</i>). If you do not need metadata, set the USE_METADATA property to false.</p>
010SK	<p>Database cannot set connection option ____.</p> <p>Description: Your application attempted an operation that the database you are connected to does not support.</p> <p>Action: You may need to upgrade your database or make sure that the latest version of metadata information is installed on it.</p>
010SL	<p>Out-of-date metadata accessor information was found on this database. Ask your database administrator to load the latest scripts.</p> <p>Description: The metadata information on the server is old and needs to be updated.</p> <p>Action: Install the stored procedures for returning metadata that come with <i>jdbc</i> (see “Installing Stored Procedures” in Chapter 3 of the <i>jdbc for JDBC Installation Guide</i>).</p>

SQL state	Message/description/action
010SM	<p>This database does not support the initial proposed set of capabilities, retrying.</p> <p>Description: Adaptive Server Enterprise versions 11.9.2 and lower had a bug that caused them to reject logins from clients that requested capabilities that the servers did not have. This warning indicates that jConnect has detected this condition and is retrying the connection using the greatest number of capabilities that the server can accept. When jConnect encounters this bug, it attempts to connect to the server twice.</p> <p>Action: Clients can safely ignore this warning, but if they wish to eliminate the warning and ensure that jConnect makes only one connection attempt, they can set the ELIMINATE_010SM connection property to <i>true</i>. Please note that this property should not be set to <i>true</i> when connecting to Adaptive Server versions 12.0 and later.</p>
010SN	<p>Permission to write to file was denied. File: _____. Error message: _____</p> <p>Description: Permission to write to a file specified in the PROTOCOL_CAPTURE connection property is denied because of a security violation in the VM. This can occur if an applet attempts to write to the specified file.</p> <p>Action: If you are attempting to write to the file from an applet, make sure that the applet has access to the target file system.</p>
010SP	<p>File could not be opened for writing. File: _____. Error message: _____</p> <p>Action: Make sure that the file name is correct and that the file is writable.</p>
010SQ	<p>The connection or login was refused, retrying connection with the host/port address.</p> <p>Description: The CONNECTION_FAILOVER connection property is set to true, and jConnect was unable to connect to one of the database servers in the list of servers to which to connect. Therefore, jConnect will now try to connect to the next server in the list.</p> <p>Action: No action is required, as long as jConnect is able to connect to another database server. However, you should determine why jConnect was unable to connect to the particular server that caused the connection warning to be issued.</p>
010TP	<p>The connection's initial character set,_____, could not be converted by the server. The server's proposed character set,_____, will be used, with conversions performed by jConnect.</p> <p>Description: The server cannot use the character set initially requested by jConnect, and has responded with a different character set. jConnect accepts the change, and will perform the necessary character-set conversions.</p> <p>The message is strictly informational and has no further consequences.</p> <p>Action: To avoid this message, set the CHARSET connection property to a character set that the server supports.</p>

SQL state	Message/description/action
010TQ	<p>jConnect could not determine the server's default character set. This is likely because of a metadata problem. Please install the required tables as mentioned in the jConnect documentation. The connection is defaulting to the <code>ascii_7</code> character set, which can handle only characters in the range from 0x00 through 0x7F.</p> <p>Description: jConnect could not determine the server's default character set. When this occurs, the only characters that are guaranteed to translate properly are the first 127 ASCII characters. Therefore, jConnect will revert to 7-bit ASCII in this case. The message is strictly informational and has no further consequences.</p> <p>Action: Install the stored procedures for returning metadata that comes with jConnect (see “Installing Stored Procedures” in Chapter 3 of the <i>jConnect for JDBC Installation Guide</i>).</p>
010UF	<p>Attempt to execute use database command failed. Error message:_____</p> <p>Description: jConnect cannot connect to the database specified in the connection URL. Two possibilities are:</p> <ul style="list-style-type: none"> • The name was entered incorrectly in the URL. • <code>USE_METADATA</code> is true (the default condition), but the stored procedures for returning metadata have not been installed. As a result, jConnect tried to execute the <code>use database</code> command with the database in the URL, but the command failed. This may be because you attempted to access an Adaptive Anywhere database. SQL Anywhere databases do not support the <code>use database</code> command. <p>Action: Make sure the database name in the URL is correct. Make sure that the stored procedures for returning metadata are installed on the server (see “Installing Stored Procedures” in Chapter 3 of the <i>jConnect for JDBC Installation Guide and Release Bulletin</i>). If you are attempting to access a SQL Anywhere database, either do not specify a database name in the URL, or set <code>USE_METADATA</code> to false.</p>
010UP	<p>Unrecognized connection property _____ ignored.</p> <p>Description: You attempted to set a connection property in the URL that jConnect does not currently recognize. jConnect ignores the unrecognized property.</p> <p>Action: Check the URL definition in your application to make sure it references only valid jConnect driver connection properties.</p>
0100V	<p>The version of TDS protocol being used is too old. Version: _____</p> <p>Description: The server does not support the required version of the TDS protocol. jConnect requires version 5.0 or later.</p> <p>Action: Use a server that supports the required version of TDS. See the jConnect installation guide's system requirements section for details.</p>
JZ001	<p>User name property '_____' too long. Maximum length is 30.</p> <p>Action: Do not exceed the 30 byte maximum.</p>
JZ002	<p>Password property '_____' too long. Maximum length is 30.</p> <p>Action: Do not exceed the 30-byte maximum.</p>

SQL state	Message/description/action
JZ003	<p>Incorrect URL format. URL: _____</p> <p>Action: Verify the URL format. See “URL connection property parameters” on page 21.</p> <p>If you are using the PROXY connection property, you may get a JZ003 exception while trying to connect if the format for the PROXY property is incorrect.</p> <p>The PROXY format for the Cascade proxy is:</p> <p><i>ip_address:port_number</i></p> <p>The PROXY format for the TDS tunnelling servlet is:</p> <p><i>http[s]://host:port/tunneling_servlet_alias</i></p>
JZ004	<p>User name property missing in DriverManager.getConnection(..., Properties)</p> <p>Action: Provide the required user property.</p>
JZ006	<p>Caught IOException: _____</p> <p>Description: An unexpected I/O error was detected from a lower layer. When such I/O exceptions are caught, they are rethrown as SQL exceptions using the ERR_IO_EXCEPTION JZ006 sqlstate. These errors are often the result of network communication problems. If the I/O exception causes the database connection to be closed, jConnect will chain a JZ0C1 exception to the JZ006. Client applications can look for the JZ0C1 exception in the chain to see if the connection is still usable.</p> <p>Action: Examine the text of the original I/O exception message, and proceed from there.</p>
JZ008	<p>Invalid column index value _____.</p> <p>Description: You have requested a column index value of less than 1 or greater than the highest available value.</p> <p>Action: Check call to the getXXX() method and the text of the original query, or be sure to call rs.next().</p>
JZ009	<p>Error encountered in conversion. Error message: _____</p> <p>Description: Some of the possibilities are:</p> <ul style="list-style-type: none"> • A conversion between two incompatible datatypes was attempted, such as date to int. • There was an attempt to convert a string containing a nonnumeric character to a numeric type. • There was a formatting error, such as an incorrectly formatted time/date string. <p>Action: Make sure that the JDBC specification supports the attempted type conversion. Make sure that strings are correctly formatted. If a string contains non-numeric characters, do not attempt to convert it to a numeric type.</p>
JZ00B	<p>Numeric overflow.</p> <p>Description: You tried to send a BigInteger as a TDS numeric, and the value was too large, or you tried to send a Java long as an int and the value was too large.</p> <p>Action: These values cannot be stored in Sybase. For long, consider using a Sybase numeric. There is no solution for Bignum.</p>

SQL state	Message/description/action
JZ00C	<p>The precision and scale specified cannot accommodate numeric value _____.</p> <p>Description: When using the <code>setBigDecimal</code> method, the <code>BigDecimal</code> value has a precision or scale that exceeds the specified precision or scale.</p> <p>Action: Make sure that the specified precision and scale can accommodate the <code>BigDecimal</code> value.</p>
JZ00E	<p>Attempt to call <code>execute()</code> or <code>executeUpdate()</code> for a statement where <code>setCursorName()</code> has been called.</p> <p>Action: Do not try to call <code>execute</code> or <code>executeUpdate</code> on a statement that has a cursor name set. Use a separate statement to delete or update a cursor. See “Using cursors with result sets” on page 48 for more information</p>
JZ00F	<p>Cursor name has already been set by <code>setCursorName()</code>.</p> <p>Action: Do not set the cursor name twice for a statement. Close the result set of the current cursor statement.</p>
JZ00G	<p>No column values were set for this row update.</p> <p>Description: You attempted to update a row in which no column values were changed.</p> <p>Action: To change column values in a row, call <code>updateXX()</code> methods before calling <code>updateRow()</code>.</p>
JZ00H	<p>The result set is not updatable. Use <code>Statement.setResultSetConcurrencyType()</code>.</p> <p>Action: To change a result set from read-only to updatable, use the <code>Statement.setResultSetConcurrencyType()</code> method or add a <code>for update</code> clause to your SQL select statement.</p>
JZ00I	<p>Invalid scale. The specified scale must be ≥ 0.</p> <p>Description: The scale value must be greater than zero.</p> <p>Action: Be sure the scale value is not negative.</p>
JZ00L	<p>Login failed. Examine the <code>SQLWarnings</code> chained to this exception for the reason(s).</p> <p>Action: See message text; proceed according to the reason(s) given for the login failure.</p>
JZ00M	<p>Login timed out. Check that your database server is running on the host and port number you specified. Also check the database server for other conditions (such as a full tempdb) that might be causing it to hang.</p> <p>Action: Follow the recommended actions in the error message.</p>
JZ010	<p>Unable to deserialize an Object value. Error text: _____</p> <p>Action: Make sure that the Java object from the database implements the <code>Serializable</code> interface and is in your local <code>CLASSPATH</code> variable.</p>

SQL state	Message/description/action
JZ011	<p>Number format exception encountered while parsing numeric connection property ____.</p> <p>Description: A noninteger value was specified for a numeric connection property.</p> <p>Action: Specify an integer value for the connection property.</p>
JZ012	<p>Internal Error. Please report it to Sybase technical support. Wrong access type for connection property ____.</p> <p>Action: Please contact Sybase Technical Support.</p>
JZ013	<p>Error obtaining JNDI entry: ____</p> <p>Action: Correct the JNDI URL or make a new entry in the directory service.</p>
JZ014	<p>You may not <code>setTransactionIsolation(Connection.TRANSACTION_NONE)</code>. This level cannot be set; it can only be returned by a server.</p> <p>Action: Check your application code, where it calls <code>Connection.setTransactionIsolation()</code>, and verify the value you are passing to the method.</p>
JZ0BD	<p>Out of range or invalid value used for method parameter.</p> <p>Action: Verify that the parameter value in the method is correct.</p>
JZ0BI	<p>Message: <code>setFetchSize</code>: The fetch size should be set with the following restrictions - <code>0 <= rows <= (maximum number of rows in the ResultSet)</code>.</p> <p>Description: The client application has tried to call <code>setFetchSize</code> with an invalid number of rows.</p> <p>Action: Verify that you are calling <code>setFetchSize</code> with the parameter falling within the above range of values.</p>
JZ0BP	<p>Output parameters are not allowed in Batch Update Statements.</p> <p>Action: Examine your application code and check that you did not try to declare an output parameter in your batch.</p>
JZ0BR	<p>The cursor is not positioned on a row that supports the ____ method.</p> <p>Description: You attempted to call a <code>ResultSet</code> method that is invalid for the current row position (for example, calling <code>insertRow()</code> when the cursor is not on the insert row).</p> <p>Action: Do not attempt to call a <code>ResultSet</code> method that is invalid for the current row position.</p>
JZ0BS	<p>Batch Statements not supported.</p> <p>Action: Install or update the <code>jConnect</code> metadata stored procedures on your database with the latest versions.</p>
JZ0BT	<p>The ____ method is not supported for <code>ResultSet</code>s of type ____.</p> <p>Description: You attempted to call a <code>ResultSet</code> method that is invalid for the type of <code>ResultSet</code>.</p> <p>Action: Do not attempt to call a <code>ResultSet</code> method that is invalid for the type of <code>ResultSet</code>.</p>

SQL state	Message/description/action
JZ0C0	<p>Connection is already closed.</p> <p>Description: The application has already called <code>Connection.close()</code> on this connection object; it cannot be used any more.</p> <p>Action: Fix the code so that connection object references are nulled out when a connection is closed.</p>
JZ0C1	<p>An <code>IOException</code> occurred which closed the connection.</p> <p>Description: An unrecoverable <code>IOException</code> occurred which caused the <code>Connection</code> to be closed. The <code>Connection</code> cannot be used for any further database activity. If this exception occurs, it will always be found in an exception chain with the JZ006 Exception (explained earlier).</p> <p>Action: Determine the cause of the <code>IOException</code> that disrupted the <code>Connection</code>.</p>
JZ0CL	<p>You must define the <code>CLASS_LOADER</code> property when using the <code>PRELOAD_JARS</code> property.</p> <p>Action: Make sure to specify a <code>CLASS_LOADER</code> when setting <code>PRELOAD_JARS</code> to a non-null value.</p>
JZ0CU	<p><code>getUpdateCount</code> can only be called once after a successful call to <code>getMoreResults</code>, or <code>execute methods</code>.</p> <p>Description: As per the JDBC API, <code>getUpdateCount</code> should be called only once per result.</p> <p>Action: Make sure your code does not call <code>getUpdateCount</code> more than once per result.</p>
JZ0D4	<p>Unrecognized protocol in Sybase JDBC URL:_____.</p> <p>Description: You specified a connection URL using a protocol other than TDS, which is the only protocol currently supported by <code>jConnect</code>.</p> <p>Action: Check the URL definition. If the URL specifies TDS as a subprotocol, make sure that the entry uses the following format and capitalization:</p> <p><code>jdbc:sybase:Tds:host:port</code></p> <p>If the URL specifies JNDI as a subprotocol, make sure that it starts with:</p> <p><code>jdbc:sybase:jndi:</code></p>
JZ0D5	<p>Error loading protocol _____.</p> <p>Action: Check the settings for the <code>CLASSPATH</code> system variable.</p>
JZ0D6	<p>Unrecognized version number _____ specified in <code>setVersion</code>. Choose one of the <code>SybDriver.VERSION_*</code> values, and make sure that the version of <code>jConnect</code> that you are using is at or beyond the version you specify.</p> <p>Action: See message text.</p>
JZ0D7	<p>Error loading url provider _____. Error message: _____</p> <p>Action: Check the JNDI URL to make sure it is correct.</p>
JZ0D8	<p>Error initializing url provider: _____</p> <p>Action: Check the JNDI URL to make sure it is correct.</p>

SQL state	Message/description/action
JZ0DP	<p>This statement has no metadata because it was not dynamically prepared. Set the DYNAMIC_PREPARE connection property to true to ensure use of dynamic statements.</p> <p>Action: Refer to the error message.</p>
JZ0EM	<p>End of data.</p> <p>Action: Please report this error to Sybase Technical Support.</p>
JZ0F1	<p>Sybase high-availability failover connection was requested but the companion server address is missing.</p> <p>Description: When you set the REQUEST_HA_SESSION connection property to true, you must also specify a failover server.</p> <p>Action: You can specify the secondary server using the SECONDARY_SERVER_HOSTPORT connection property, or you can set the secondary server using JNDI (see “Implementing high availability failover support” on page 40).</p>
JZ0F2	<p>Sybase high-availability failover has occurred. The current transaction is aborted, but the connection is still usable. Retry your transaction.</p> <p>Description: The backend database server to which you were connected has gone down, but you have failed over to a secondary server. The database connection is still usable.</p> <p>Action: Client code should catch this exception, then restart the transaction from the last committed point. Assuming you properly handle the exception, you can continue executing JDBC calls on the same Connection object.</p>
JZ0H0	<p>Unable to start thread for event handler; event name = ____.</p> <p>Action: Please report this error to Sybase Technical Support.</p>
JZ0H1	<p>An event notification was received but the event handler was not found; event name = ____.</p> <p>Action: Please report this error to Sybase Technical Support.</p>
JZ0HC	<p>Illegal character '____' encountered while parsing hexadecimal number.</p> <p>Description: A string that is supposed to represent a binary value contains a character that is not in the range (0–9, a–f) that is required for a hexadecimal number.</p> <p>Action: Check the character values in the string to make sure they are in the required range.</p>
JZ0I1	<p>I/O Layer: Error reading stream.</p> <p>Description: The connection was unable to read the amount requested. Most likely, the statement timeout period was exceeded and the connection timed out.</p> <p>Action: Increase the statement timeout value.</p>
JZ0I2	<p>I/O layer: Error writing stream.</p> <p>Description: The connection was unable to write the output requested. Most likely, the statement time-out period was exceeded and the connection timed out.</p> <p>Action: Increase the statement time out value.</p>

SQL state	Message/description/action
JZ013	<p>Unknown property. This message indicates an internal product problem. Report this error to Sybase Technical support.</p> <p>Action: Indicates an internal product problem. Please report this error to Sybase Technical Support.</p>
JZ015	<p>An unrecognized CHARSET property was specified: ____.</p> <p>Description: You specified an unsupported character set code for the CHARSET connection property.</p> <p>Action: Enter a valid character-set code for the connection property. See “jConnect character-set converters” on page 34.</p>
JZ016	<p>An error occurred converting UNICODE to the charset used by the server. Error message: ____</p> <p>Action: Choose a different character set code for the CHARSET connection property on the jConnect client that can support all the characters you need to send to the server. You may need to install a different character set on the server, too. Also, if you are using jConnect version 5.5 or later, and Adaptive Server Enterprise 12.5 or later, you can send your data to the server as unichar/univarchar datatypes. Please see “Using jConnect to pass Unicode data” on page 33.</p>
JZ017	<p>No response from proxy gateway.</p> <p>Description: The Cascade or security gateway is not responding.</p> <p>Action: Make sure the gateway is properly installed and running.</p>
JZ018	<p>Proxy gateway connection refused. Gateway response: ____</p> <p>Description: The Web server/gateway indicated by the PROXY connection property has refused your connection request.</p> <p>Action: Check the access and error logs on the proxy to determine why the connection was refused. Make sure the proxy is a JDBC gateway.</p>
JZ019	<p>This InputStream was closed.</p> <p>Description: You tried to read an InputStream obtained from getAsciiStream(), getUnicodeStream(), or getBinaryStream(), but the InputStream was already closed. The stream may have been closed because you moved to another column or cancelled the result set and there were not enough resources to cache the data.</p> <p>Action: Increase the cache size or read columns in order.</p>
JZ01A	<p>Truncation error trying to send ____.</p> <p>Description: There was a truncation error on character set conversion prior to sending a string. The converted string is longer than the size allocated for it.</p> <p>Action: Choose a different character-set code for the CHARSET connection property on the jConnect client that can support all the characters you need to send to the server. You may need to install a different character set on the server, too.</p>
JZ01R	<p>getXXX may not be called on a column after it has been updated in the result set with a java.io.Reader.</p> <p>Action: Remove the getXXX call on the ResultSet column which you updated using a Reader.</p>

SQL state	Message/description/action
JZ0IS	<p>getXXXStream may not be called on a column after it has been updated in the result set.</p> <p>Description: After updating a column in a result set, you attempted to read the updated column value using one of the following SybResultSet methods: getAsciiStream(), getUnicodeStream(), getBinaryStream(). jConnect does not support this usage.</p> <p>Action: Do not attempt to fetch input streams from columns you are updating.</p>
JZ0JO	<p>Offset and/or length values exceed the actual text/image length.</p> <p>Action: Check the offset and/or length values you used to make sure they are correct.</p>
JZ0LC	<p>You cannot call the ____ method on a ResultSet which is using a language cursor to fetch rows. Try setting the LANGUAGE_CURSOR connection property to false.</p> <p>Description: The application tried to call one of the ResultSet cursor scrolling methods on a ResultSet which was created with a language cursor.</p> <p>Action: See the error message.</p>
JZ0NC	<p>wasNull called without a preceding call to get a column.</p> <p>Description: You can only call wasNull() after a call to get a column, such as getInt() or getBinaryStream().</p> <p>Action: Change the code to move the call to wasNull().</p>
JZ0NE	<p>Incorrect URL format. URL: _____. Error message: _____</p> <p>Action: Check the format of the URL. Make sure that the port number consists only of numeric characters.</p>
JZ0NF	<p>Unable to load SybSocketFactory. Make sure that you have spelled the class name correctly, that the package is fully specified, that the class is available in your class path, and that it has a public zero-argument constructor.</p> <p>Action: See message text.</p>
JZ0P1	<p>Unexpected result type.</p> <p>Description: The database has returned a result that the statement cannot return to the application, or that the application is not expecting at this point. This generally indicates that the application is using JDBC incorrectly to execute the query or stored procedure. If the JDBC application is connected to an Open Server application, it may indicate an error in the Open Server application that causes the Open Server to send unexpected sequences of results.</p> <p>Action: Use the com.sybase.utils.Debug(true, "ALL") debugging tools to try to figure out what unexpected result is seen, and to understand its causes.</p>
JZ0P4	<p>Protocol error. This message indicates an internal product problem. Report this error to Sybase technical support.</p> <p>Action: See message text.</p>

SQL state	Message/description/action
JZ0P7	<p>Column is not cached; use RE-READABLE_COLUMNS property.</p> <p>Description: With the REPEAT_READ connection property set to false, an attempt was made to reread a column or read a column in the wrong order.</p> <p>When REPEAT_READ is false, you can only read the column value for a row once, and you can only read columns in ascending column-index order. For example, after reading column 3 for a row, you cannot read its value a second time and you cannot read column 2 for the row.</p> <p>Action: Either set REPEAT_READ to true, or do not attempt to reread a column value and make sure that you read columns in ascending column-index order.</p>
JZ0P8	<p>The RSMDBA Column Type Name you requested is unknown.</p> <p>Description: jConnect cannot determine the name of a column type in the ResultSetMetaData.getColumnTypeName() method.</p> <p>Action: Make sure that your database has the most recent stored procedures for metadata.</p>
JZ0P9	<p>A COMPUTE BY query has been detected. That type of result is unsupported and has been cancelled.</p> <p>Description: The query you executed returned COMPUTE results, which are not supported by jConnect.</p> <p>Action: Change your query or stored procedure so it does not use COMPUTE BY.</p>
JZ0PA	<p>The query has been cancelled and the same response discarded.</p> <p>Description: A cancel was probably issued by another statement on the connection.</p> <p>Action: Check the chain of SQL exceptions and warnings on this and other statements to determine the cause.</p>
JZ0PB	<p>The server does not support a requested operation.</p> <p>Description: When jConnect creates a connection with a server, it informs the server of capabilities it wants supported and the server informs jConnect of the capabilities that it supports. This error message is sent when an application requests an operation that was denied in the original capabilities negotiation.</p> <p>For example, if the database does not support precompilation of dynamic SQL statements, and your code invokes SybConnection.prepareStatement(<i>sql_stmt</i>, <i>dynamic</i>), and <i>dynamic</i> is set to true, jConnect generates this message.</p> <p>Action: Modify your code so that it does not request an unsupported capability.</p>

SQL state	Message/description/action
JZ0PC	<p>The number and size of parameters in your query require wide table support. But either the server does not offer such support, or it was not requested during the login sequence. Try setting the JCONNECT_VERSION property to <code>>=6</code> if you wish to request widetable support.</p> <p>Description: You are trying to execute a statement with a large number of parameters, and the server is not configured to handle that many parameters. The number of parameters that can produce this exception will vary, depending on the datatypes of the data you are sending. You will never get this exception if you are sending 481 or fewer parameters.</p> <p>Action: You must run this query against an Adaptive Server 12.5 or later server. When you connect to the database, set the JCONNECT_VERSION property to <code>VERSION_6</code> or <code>VERSION_LATEST</code>.</p>
JZ0PD	<p>The size of the query in your dynamic prepare is large enough that you require widetable support. But either the server does not offer such support, or it was not requested during the login sequence. Try setting the JCONNECT_VERSION property to <code>>=6</code> if you wish to request widetable support.</p> <p>Description: You are trying to execute a dynamic prepared statement with a large number of parameters, and the server is not configured to handle that many parameters.</p> <p>Action: You must run this query against an Adaptive Server 12.5 or later server. When you connect to the database, set the JCONNECT_VERSION property to <code>VERSION_6</code> or <code>VERSION_LATEST</code>.</p>
JZ0PE	<p>The number of columns in your cursor declaration OR the size of your cursor declaration itself are large enough that you require widetable support. But either the server does not offer such support, or it was not requested during the login sequence. Try setting the JCONNECT_VERSION property to <code>>= 6</code> if you wish to request wide table support.</p> <p>Description: This error can occur when your SELECT statement tries to return data from more than 255 columns, or when the actual length of the SELECT statement is very large (greater than approximately 65500 characters).</p> <p>Action: You must run this query against a version 12.5 or later Adaptive Server. When you connect to the database, set the JCONNECT_VERSION property to <code>VERSION_6</code> or <code>VERSION_LATEST</code>.</p>
JZ0PN	<p>Specified port number of ____ was out of range. Port numbers must meet the following conditions: <code>0<= portNumber <=65535</code>.</p> <p>Action: Check the port number that is specified in the database URL.</p>

SQL state	Message/description/action
JZOR0	<p>Result set has already been closed.</p> <p>Description: The <code>ResultSet.close()</code> method has already been called on the result set object; you cannot use the result set for anything else.</p> <p>Action: Fix the code so that <code>ResultSet</code> object references are set to null whenever a result set is closed.</p>
JZOR1	<p>Result set is IDLE as you are not currently accessing a row.</p> <p>Description: The application has called one of the <code>ResultSet.getXXX</code> column-data retrieval methods, but there is no current row; the application has not called <code>ResultSet.next()</code>, or <code>ResultSet.next()</code> returned false to indicate that there is no data.</p> <p>Action: Check that <code>rs.next()</code> is set to true before calling <code>rs.getXXX</code>.</p>
JZOR2	<p>No result set for this query.</p> <p>Description: You used <code>Statement.executeQuery()</code>, but the statement returned no rows.</p> <p>Action: Use <code>executeUpdate</code> for statements returning no rows.</p>
JZOR3	<p>Column is DEAD. This is an internal error. Please report it to Sybase technical support.</p> <p>Action: See message text.</p>
JZOR4	<p>Column does not have a text pointer. It is not a text/image column or the column is NULL.</p> <p>Description: You cannot update a text/image column if it is NULL. A NULL text/image column does not contain a text pointer.</p> <p>Action: Make sure that you are not trying to update or get a text pointer to a column that does not support text/image data. Make sure that you are not trying to update a text/image column that is NULL. Insert data first, then make the update.</p>
JZOR5	<p>The <code>ResultSet</code> is currently positioned beyond the last row. You cannot perform a <code>get*</code> operation to read data in this state.</p> <p>Description: The application has moved the <code>ResultSet</code> row pointer beyond the last row. In this position, there is no data to read, so any <code>get*</code> operations are illegal.</p> <p>Action: Alter the code so that it does not attempt to read column data when the <code>ResultSet</code> is positioned beyond the last row.</p>
JZORD	<p>You cannot call any of the <code>ResultSet.get*</code> methods on a row that has been deleted with the <code>deleteRow()</code> method.</p> <p>Description: An application is trying to retrieve data from a row that has been deleted. There is no valid data to be retrieved.</p> <p>Action: Alter the code so that the application does not attempt to retrieve data from a deleted row.</p>

SQL state	Message/description/action
JZ0RM	<p>refreshRow may not be called after updateRow or deleteRow.</p> <p>Description: After updating a row in the database with SybCursorResult.updateRow(), or deleting it with SybCursorResult.deleteRow(), you used SybCursorResult.refreshRow() to refresh the row from the database.</p> <p>Action: Do not attempt to refresh a row after updating it or deleting it from the database.</p>
JZ0S0	<p>Statement state machine: Statement is BUSY.</p> <p>Description: The only time this error is raised is from the Statement.setCursorname() method, if the application is trying to set the cursor name when the statement is already in use and has noncursor results that need to be read.</p> <p>Action: Set the cursor name on a statement before you execute any queries on it, or call Statement.cancel() before setting the cursor name, to make sure that the statement isn't busy.</p>
JZ0S1	<p>Statement state machine: Trying to FETCH on IDLE statement.</p> <p>Description: An internal error occurred on the statement.</p> <p>Action: Close the statement and open another one.</p>
JZ0S2	<p>Statement object has already been closed.</p> <p>Description: The Statement.close() method has already been called on the statement object; you cannot use the statement for anything else.</p> <p>Action: Fix the application so that statement object references are set to null whenever a statement is closed.</p>
JZ0S3	<p>The inherited method _____ cannot be used in this subclass.</p> <p>Description: PreparedStatement does not support executeQuery(String), executeUpdate(String), or execute(String).</p> <p>Action: To pass a query string, use Statement, not PreparedStatement.</p>
JZ0S4	<p>Cannot execute an empty (zero-length) query.</p> <p>Action: Do not execute an empty query ("").</p>
JZ0S5	<p>The local transaction method _____ cannot be used while a global transaction is active on this connection.</p> <p>Description: This exception can occur when using distributed transactions.</p> <p>Action: See Chapter 7, "Distributed Transactions," in the <i>JDBC 2.0 Optional Package</i> (formerly the JDBC 2.0 Standard Extension API) for more information on diagnosing the problem.</p>
JZ0S6	<p>The local transaction method _____ cannot be used on a pre-System 12 XAConnection.</p> <p>Description: This exception can occur when using distributed transactions.</p> <p>Action: See Chapter 7, "Distributed Transactions," in the <i>JDBC 2.0 Optional Package</i> (formerly the JDBC 2.0 Standard Extension API) for more information on diagnosing the problem.</p>
JZ0S8	<p>An escape sequence in a SQL Query was malformed: '_____'.</p> <p>Description: This error results from bad escape syntax.</p> <p>Action: Check JDBC documentation for correct syntax.</p>

SQL state	Message/description/action
JZ0S9	<p>Cannot execute an empty (zero-length) query.</p> <p>Action: Do not execute an empty query ("").</p>
JZ0SA	<p>Prepared Statement: Input parameter not set, index: ____.</p> <p>Action: Make sure that each input parameter has a value.</p>
JZ0SB	<p>Parameter index out of range: ____.</p> <p>Description: You have attempted to get, set, or register a parameter that goes beyond the maximum number of parameters.</p> <p>Action: Check the number of parameters in your query.</p>
JZ0SC	<p>Callable Statement: attempt to set the return status as an InParameter.</p> <p>Description: You have prepared a call to a stored procedure that returns a status, but you are trying to set parameter 1, which is the return status.</p> <p>Action: Parameters that you can set start at 2 with this type of call.</p>
JZ0SD	<p>No registered parameter found for output parameter.</p> <p>Description: This indicates an application logic error. You attempted to call getXXX() or wasNull() on a parameter, but you have not read any parameters yet, or there are no output parameters.</p> <p>Action: Check to make sure that the application has registered output parameters on the CallableStatement, that the statement has been executed, and that the output parameters were read.</p>
JZ0SE	<p>Invalid object type specified for setObject().</p> <p>Description: Illegal type argument passed to PreparedStatement.setObject.</p> <p>Action: Check the JDBC documentation. The argument must be a constant from java.sql.Types.</p>
JZ0SF	<p>No Parameters expected. Has query been sent?</p> <p>Description: You tried to set a parameter on a statement with no parameters.</p> <p>Action: Make sure the query has been sent before you set the parameters.</p>
JZ0SG	<p>An RPC did not return as many output parameters as the application had registered for it.</p> <p>Description: This error occurs if you call CallableStatement.registerOutParam() for more parameters than you declared as "OUTPUT" parameters in the stored procedure. See "RPC returns fewer output parameters than registered" on page 108 for more information.</p> <p>Action: Check your stored procedures and registerOutParameter calls. Make sure that you have declared all of the appropriate parameters as "OUTPUT." Look at the line of code that reads:</p> <pre>create procedure yourproc (@p1 int OUTPUT, ...</pre> <p>Note If you receive this error while using Adaptive Server Anywhere (previously known as SQL Anywhere), upgrade to Adaptive Server Anywhere version 5.5.04.</p>

SQL state	Message/description/action
JZ0SH	<p>A static function escape was used, but the metadata accessor information was not found on this server.</p> <p>Action: Install metadata accessor information before using static function escapes.</p>
JZ0SI	<p>A static function escape _____ was used which is not supported by this server.</p> <p>Action: Do not use this escape.</p>
JZ0SJ	<p>Metadata accessor information was not found on this database.</p> <p>Action: Install metadata information before making metadata calls.</p>
JZ0SK	<p>The oj escape is not supported for this type of database server. Workaround: use server-specific outer join syntax, if supported. Consult server documentation.</p> <p>Action: Read the error message. Also, install the latest version of the jConnect metadata.</p>
JZ0SL	<p>Unsupported SQL type _____.</p> <p>Description: The application has declared a parameter to be of a type that jConnect does not support.</p> <p>Action: If possible, try declaring the parameter to be of a different type. Do not use Types.NULL or PreparedStatement.setObject (null).</p>
JZ0SM	<p>jConnect could not execute a stored procedure because there was a problem sending the parameter(s). This problem was likely caused because the server does not support a specific datatype, or because jConnect did not request support for that datatype at connect time. Try setting the JCONNECT_VERSION connection property to a higher value. Or, if possible, try sending your procedure execution command as a language statement.</p>
JZ0SN	<p>setMaxFieldSize: field size cannot be negative.</p> <p>Action: Use a positive value or zero (unlimited) when calling setMaxFieldSize.</p>
JZ0SO	<p>Invalid ResultSet concurrency type:_____.</p> <p>Action: Check that your declared concurrency is either ResultSet.CONCUR_READ_ONLY or ResultSet.CONCUR_UPDATABLE.</p>
JZ0SP	<p>Invalid ResultSet type:_____.</p> <p>Action: Check that your declared ResultSet type is ResultSet.TYPE_FORWARD_ONLY or ResultSet.TYPE_SCROLL_INSENSITIVE. jConnect does not support the ResultSet.TYPE_SCROLL_SENSITIVE ResultSet type.</p>
JZ0SQ	<p>In valid UDT type _____.</p> <p>Description: When calling the DatabaseMetaData.getUDTs method, jConnect will throw this Exception if the user-defined type is not either Types.JAVA_OBJECT, Types.STRUCT or Types.DISTINCT.</p> <p>Action: Use one of the three UDTs mentioned above.</p>

SQL state	Message/description/action
JZ0SR	<p>setMaxRows: max rows cannot be negative.</p> <p>Action: Use a positive value or zero (unlimited) when calling setMaxRows.</p>
JZ0SS	<p>setQueryTimeout: query timeout cannot be negative.</p>
JZ0ST	<p>jConnect cannot send a Java object as a literal parameter in a query. Make sure that your database server supports Java objects and that the LITERAL_PARAMS connection property is set to false when you execute this query.</p>
JZ0SU	<p>A Date or Timestamp parameter was set with a year of _____, but the server can only support year values between _____ and _____. If you're trying to send data to date or timestamp columns or parameters on Adaptive Server Anywhere, you may wish to send your data as Strings, and let the server convert them.</p> <p>Description: Adaptive Server Enterprise and Adaptive Server Anywhere have different allowable ranges for datetime and date values. Datetime values must have years greater or equal to 1753. The date datatype, however, can hold years greater or equal to 1.</p> <p>Action: Make sure that the date/timestamp value you are sending falls in the acceptable range.</p>
JZ0T2	<p>Listener thread read error.</p> <p>Action: Check your network communications.</p>
JZ0T3	<p>Read operation timed out.</p> <p>Description: The time allotted to read the response to a query was exceeded.</p> <p>Action: Increase the timeout period by calling Statement.setQueryTimeout().</p>
JZ0T4	<p>Write operation timed out. Timeout in milliseconds: _____.</p> <p>Description: The time allotted to send a request was exceeded.</p> <p>Action: Increase the timeout period by calling Statement.setQueryTimeout().</p>
JZ0T5	<p>Cache used to store responses is full.</p> <p>Action: Use default or larger value for the STREAM_CACHE_SIZE connection property.</p>
JZ0T6	<p>Error reading tunneled TDS URL.</p> <p>Description: The tunneled protocol failed while reading the URL header.</p> <p>Action: Check the URL you defined for the connection.</p>
JZ0T7	<p>Listener thread read error -- caught ThreadDeath. Check network connection.</p> <p>Action: Check the network connections and try to run the application again. If the threads continue to be aborted, please contact Sybase Technical Support.</p>
JZ0T8	<p>Data received for an unknown request. Please report this error to Sybase Technical Support.</p>
JZ0T9	<p>Request to send not synchronized. Please report this error to Sybase Technical Support.</p> <p>Action: See message text.</p>

SQL state	Message/description/action
JZ0TC	<p>Attempted conversion between an illegal pair of types.</p> <p>Description: Conversion between a Java type and a SQL type failed.</p> <p>Action: Check the requested type conversion to make sure it is supported in the JDBC specification.</p>
JZ0TE	<p>Attempted conversion between an illegal pair of types. Valid database types are: '_____.'</p> <p>Description: The database column datatype and the datatype requested in the <code>ResultSet.getXXX()</code> call are not implicitly convertible.</p> <p>Action: Use one of the valid datatypes listed in the error message.</p>
JZ0TI	<p><code>jConnect</code> cannot make a meaningful conversion between the database type of _____ and the requested type of _____.</p> <p>Description: This kind of exception can occur, for example, if an application tries to call <code>ResultSet.getObject(int, Types.DATE)</code> on a time value that is returned from the database.</p> <p>Action: Make sure that the database datatype is implicitly convertible to the Object type you wish to retrieve.</p>
JZ0TO	<p>Read operation timed out.</p> <p>Description: This exception occurs when there is a socket read timeout.</p> <p>Action: Increase the timeout period by calling <code>Statement.setQueryTimeout()</code>. Also, check the query or stored proc you are executing to determine why it is taking longer than expected.</p>
JZ0TS	<p>Truncation error trying to send _____.</p> <p>Description: The application specified a String that was longer than the length that the application wanted to send. Therefore, the String is truncated to the declared length.</p> <p>Action: Set the length properly to avoid truncation.</p>
JZ0US	<p>The <code>SybSocketFactory</code> connection property was set, and the <code>PROXY</code> connection property was set to the URL of a servlet. The <code>jConnect</code> driver does not support this combination. If you want to send secure HTTP from an applet running within a browser, use a proxy URL beginning with "https://".</p> <p>Action: See message text.</p>
JZ0XC	<p>_____ is an unrecognized transaction coordinator type.</p> <p>Description: The metadata information indicates that the server supports distributed transactions, but <code>jConnect</code> does not support the protocol being used.</p> <p>Action: Verify that you have installed the latest metadata scripts. If the error persists, please contact Sybase Technical Support.</p>

SQL state	Message/description/action
JZ0XS	<p>The server does not support XA-style transactions. Please verify that the transaction feature is enabled and licensed on this server.</p> <p>Description: The server to which jConnect attempted a connection does not support distributed transactions.</p> <p>Action: Do not use XADatasource with this server, or upgrade or configure the server for distributed transactions.</p>
JZ0XU	<p>Current user does not have permission to do XA-style transactions. Be sure user has _____ role.</p> <p>Description: The user connected to the database is not authorized to conduct distributed transactions. Most likely because they do not have the proper role (shown in the blank).</p> <p>Action: Grant the user the role shown in the error message, or have another user with that role conduct the transaction.</p>
S0022	<p>Invalid column name '_____'. </p> <p>Description: You attempted to reference a column by name and there is no column with that name.</p> <p>Action: Check the spelling of the column name.</p>
ZZ00A	<p>The method _____ has not been completed and should not be called.</p> <p>Description: You attempted to use a method that is not implemented.</p> <p>Action: Check the release bulletin that came with your version of jConnect for further information. You can also check the jConnect Web page at http://www.sybase.com to see whether a more recent version of jConnect implements the method. If not, do not use the method.</p>

This appendix is a guide to jConnect sample programs.

Topics

Running IsqlApp

Running jConnect sample programs and code

Running IsqlApp

IsqlApp allows you to issue isql commands from the command line, and run jConnect sample programs.

The syntax for IsqlApp is:

```
IsqlApp [-U username] [-P password]
        [-S servername]
        [-G gateway]
        [-p {http|https}]
        [-D debug-class-list]
        [-v]
        [-I input-command-file]
        [-c command_terminator]
        [-C charset] [-L language]
        [-T sessionID]
        [-V <version {2,3,4,5}>]
```

Parameter	Description
-U	The login ID with which you want to connect to a server.
-P	The password for the specified login ID.
-S	The name of the server to which you want to connect.
-G	Gateway address. For the HTTP protocol, the URL is: <code>http://host:port</code> . To use the HTTPS protocol that supports encryption, the URL is <code>https://host:port/servlet_alias</code> .
-p	Specifies whether you want to use the HTTP protocol or the HTTPS protocol that supports encryption.

Parameter	Description
-D	Turns on debugging for all classes or for just the ones you specify, separated by a comma. For example, <i>-D ALL</i> displays debugging output for all classes. <i>-D SybConnection, Tds</i> displays debugging output only for the SybConnection and Tds classes.
-v	Turns on verbose output for display or printing.
-I	Causes IsqlApp to take commands from a file instead of the keyboard. After the parameter, you specify the name of the file to use for the IsqlApp input. The file must contain command terminators (“go” by default).
-c	Lets you specify a keyword (for example, “go”) that, when entered on a line by itself, terminates the command. This lets you enter multiline commands before using the terminator keyword. If you do not specify a command terminator, each new line terminates a command.
-C	Specifies the character set for strings passed through TDS. If you do not specify a character set, IsqlApp uses the server’s default charset.
-L	The language in which to display error messages returned from the server and for jConnect messages.
-T	When this parameter is set, jConnect assumes that an application is trying to resume communication on an existing TDS session held open by the TDS-tunnelling gateway. jConnect skips the login negotiations and forwards all requests from the application to the specified session ID.
-V	Enables the use version-specific characteristics. See “JCONNECT_VERSION connection property” on page 10.

Note You must enter a space after each option flag.

To obtain a full description of the command-line options, enter:

```
java IsqlApp -help
```

The following example shows how to connect to a database on a host named “myserver” through port “3756” and run an isql script named “myscript”:

```
java IsqlApp -U sa -P sapassword  
-S jdbc:sybase:Tds:myserver:3756
```

```
-I $JDBC_HOME/sp/myscript -c run
```

Note An applet that provides GUI access to isql commands is available as:

`$JDBC_HOME/sample2/gateway.html` (UNIX)
`%JDBC_HOME%\sample2\gateway.html` (Windows)

Running jConnect sample programs and code

jConnect includes several sample programs that illustrate many of the topics covered in this chapter, and to help you understand how jConnect works with various JDBC classes and methods. In addition, this section includes a sample code fragment for your reference.

Sample applications

When you install jConnect, you can also install sample programs. These samples include the source code so that you can review how jConnect implements various JDBC classes and methods. See the *jConnect for JDBC Installation Guide* for complete instructions for installing the sample programs.

Note The jConnect sample programs are intended for demonstration purposes only.

The sample programs are installed in the *sample2* subdirectory under your jConnect installation directory. The file *index.html* in the *sample2* subdirectory contains a complete list of the samples that are available along with a description of each sample. *index.html* also lets you view and run the sample programs as applets.

Running the sample applets

Using your Web browser, you can run some of the sample programs as applets. This enables you to view the source code while viewing the output results.

To run the samples as applets, you need to start the Web server gateway.

Use your Web browser to open *index.html*:

Enter:

http://localhost:8000/sample2/index.html

Running the sample programs with Adaptive Server Anywhere

All of the sample programs are compatible with Adaptive Server, but only a limited number are compatible with Adaptive Server Anywhere. Refer to *index.html* in the *sample2* subdirectory for a current list of the sample programs that are compatible with Adaptive Server Anywhere.

To run the sample programs that are available for Adaptive Server Anywhere, you must install the *pubs2_any.sql* script on your Adaptive Server Anywhere server. This script is located in the *sample2* subdirectory.

For Windows, go to DOS command window and enter:

```
java IsqlApp -U dba -P password
-S jdbc:sybase:Tds:[hostname]:[port]
-I %JDBC_HOME%\sample2\pubs2_any.sql -c go
```

For UNIX, enter:

```
java IsqlApp -U dba -P password
-S jdbc:sybase:Tds:[hostname]:[port]
-I $JDBC_HOME/sample2/pub2_any.sql -c go
```

Sample code

The following sample code illustrates how to invoke the jConnect driver, make a connection, issue a SQL statement, and process results.

```
import java.io.*;
import java.sql.*;

public class SampleCode
{
    public static void main(String args[])
    {
        try
        {
            /*
             * Open the connection. May throw a SQLException.
             */
        }
    }
}
```



```
        DriverManager.registerDriver(
            (Driver) Class.forName(
                "com.sybase.jdbc2.jdbc.SybDriver").newInstance());

        Connection con = DriverManager.getConnection(
            "jdbc:sybase:Tds:myserver:3767", "sa", "");
    /*
     * Create a statement object, the container for the SQL
     * statement. May throw a SQLException.
     */
        Statement stmt = con.createStatement();
    /*
     * Create a result set object by executing the query.
     * May throw a SQLException.
     */
        ResultSet rs = stmt.executeQuery("Select 1");
    /*
     * Process the result set.
     */

        if (rs.next())
        {
            int value = rs.getInt(1);
            System.out.println("Fetched value " + value);
        }
        rs.close()
        stmt.close()
        con.close()
    } //end try
/*
 * Exception handling.
 */
    catch (SQLException sqe)
    {
        System.out.println("Unexpected exception : " +
            sqe.toString() + ", sqlstate = " +
            sqe.getSQLState());
        System.exit(1);
    } //end catch

    catch (Exception e)
    {
        e.printStackTrace();
        System.exit(1);
    } //end catch
```

```
        System.exit(0);  
    }  
}
```

Index

A

- Adaptive Server
 - connecting to 19
 - connection example 20
- Adaptive Server Anywhere 18
 - accessing metadata 48
 - connecting to 21
 - euro symbol 38
 - sending image data 63, 66
 - SERVICENAME connection property 20
 - storing and retrieving Java objects 76
- advanced features 68
- applets 132
- APPLICATIONNAME connection property 12
- applications
 - turning off debugging in 102
 - turning on debugging in 102
- ASE datatypes
 - date, time, and datetime 67
- Audience vii

B

- batch updates 61
 - stored procedures 60
- BE_AS_JDBC_COMPLIANT_ 12
- BigDecimal rescaling
 - improving driver performance 112

C

- CANCEL_ALL connection property 6, 10, 12
- capturing TDS communication 104
- character sets
 - setting 35
 - supported 36
- character-set conversion

- improving driver performance 113
- improving performance 36
- character-set converter classes
 - PureConverter 34
 - selecting 35
 - TruncationConverter 34
- character-sets
 - converter classes 34
- CHARSET connection property 6, 13
 - setting 35
- CHARSET_CONVERTER connection property 6
- CHARSET_CONVERTER_CLASS connection
 - property 13, 35
- CLASS_LOADER connection property 13
- CLASSPATH
 - setting for debugging 103
- columns
 - deletions in cursor result sets 52
 - updating in cursor result sets 53
- compute statements 98
- connecting to
 - a server using JNDI 22
 - Adaptive Server 19
 - Adaptive Server Anywhere 21
- connection
 - errors 106, 107
 - pooling 89
- Connection properties
 - EXPIRESTRING 14
- connection properties
 - APPLICATIONNAME 12
 - BE_AS_JDBC_COMPLIANT_ 12
 - CANCEL_ALL 6, 10, 12
 - CHARSET 6, 13
 - CHARSET_CONVERTER 6
 - CHARSET_CONVERTER_CLASS 13, 35
 - CLASS_LOADER 13
 - CONNECTION_FAILOVER 13, 22
 - DISABLE_UNICHAR_SENDING 13
 - DISABLE_UNPROCESSED_PARAM_WARNING

- S 14
 - DYNAMIC_PREPARE 14
 - ESCAPE_PROCESSING_DEFAULT 14, 120
 - FAKE_METADATA 14
 - GET_BY_NAME_USES_COLUMN_LABEL 15
 - HOSTNAME 15
 - HOSTPROC 15
 - IGNORE_DONE_IN_PROC 15
 - IS_CLOSED_TEST 15
 - JCONNECT_VERSION 10, 15
 - LANGUAGE 6, 15
 - LANGUAGE_CURSOR 15, 121
 - LANGUAGE_CURSOR and cursor performance 120
 - LITERAL_PARAMS 16
 - PACKETSIZE 16
 - password 16
 - PRELOAD_JARS 16
 - PROTOCOL_CAPTURE 17
 - PROXY 17
 - REMOTEPWD 17
 - REPEAT_READ 17, 112
 - REQUEST_HA_SESSION 17
 - RMNAME 17
 - SECONDARY_SERVER_HOSTPORT 18
 - SELECT_OPENS_CURSOR 18
 - SERIALIZE_REQUESTS 18
 - SERVICENAME 18
 - SESSION_ID 18
 - SESSION_TIMEOUT 19
 - setting 11
 - setting in URL 21
 - SQLINITSTRING 19
 - STREAM_CACHE_SIZE 19
 - SYB SOCKET_FACTORY 19
 - USE_METADATA 16
 - user 19
 - VERSIONSTRING 19
 - CONNECTION_FAILOVER connection property 13, 22
 - connections
 - gateway connection refused 106
 - Conventions ix
 - creating a cursor 49
 - currency symbol, euro 38
 - cursor performance 120
 - and the LANGUAGE_CURSOR connection property 120
 - cursor result sets
 - deleting a row 54
 - deletions 52
 - inserting a row 55
 - methods for updating the database 53
 - positioned updates 52
 - positioned updates and deletes using JDBC 1.x methods 51
 - positioned updates and deletes using JDBC 2.0 methods 52
 - updating columns 53
 - cursors 48
 - creating 49
 - using with a PreparedStatement 55
- D**
- data
 - image 63
 - databases
 - JNDI for naming 85
 - storing Java objects as column data in a table 76
 - datatypes
 - ASE date, time, and datetime 67
 - JDBC Date, Time, and Timestamp datatypes 67
 - unichar and univarchar 33
 - Debug class 101
 - Debug servlet argument 136
 - debugging 101
 - methods 103
 - obtaining an instance of the Debug class 101
 - setting CLASSPATH 103
 - turning off in your application 102
 - turning on in your application 102
 - deserialization 83
 - deviations from JDBC standards 95
 - DISABLE_UNICHAR_SENDING connection property 13
 - DISABLE_UNPROCESSED_PARAM_WARNINGS
 - connection property 14
 - distributed transaction support 91
 - driver
 - JDBC types 2
 - properties 12
 - dynamic class loading 80

DYNAMIC_PREPARE connection property 14

E

error messages

- customizing handling 73
- error-message handler example 74
- handling 71
- installing an error-message-handler 74
- SQL exception and warning 141
- Sybase-specific 71

errors

- connection 106, 107
- stored procedure 108

ESCAPE_PROCESSING_DEFAULT connection property 120

ESCAPE_PROCESSING_DEFAULT property 14

euro currency symbol 38

event notification 69
example 70

EXPIRESTRING connection property 14

extension changes, Sybase 125

F

FAKE_METADATA connection property 14

Font conventions ix

G

gateways 127

- configuration 128
- connection refused 106
- Open Server 22

GET_BY_NAME_USES_COLUMN_LABEL connection property 15

H

handling

- error messages 71

high availability (HA) support 40

HOSTNAME connection property 15

HOSTPROC connection property 15

HTTP 127

I

IGNORE_DONE_IN_PROC connection property 15

image data

- executing the update with TextPointer.sendData 65

getting a TextPointer object 64

public methods in the TextPointer class 63

sending 63

updating a column with TextPointer.sendData() 64

installing

- an error-message-handler 74

the TDS servlet 136

interfaces, JDBC 1

internationalization 32

invoking jConnect 11

IS_CLOSED_TEST connection property 15

Isql applet

- running the sample 133

IsqlApp utility 163

J

JAR files

- preloading 84

Java objects

- storing and retrieving in ASA 6.0 76

storing as column data in a table 76

java objects

- storing as column data in a table 76

jConnect

- debugging 101

definition 2

gateways 127

improving performance 111

invoking 11

memory problems in applications 107

sample programs 165

setting connection properties 11

setting up 5

Index

- using cursors 48
- JCONNECT_VERSION connection property 10, 15
- JDBC
 - definition 1
 - driver types 2
 - interfaces 1
 - restrictions, limitations, and deviations 95
- JDBC 2.0
 - optional package extensions support 84
 - standard extensions 84
- JDBC datatypes
 - Date, Time, and Timestamp 67
- JDBC drivers
 - JDBC-ODBC bridge 2
 - native-API/partly-Java 2
 - native-protocol/all-Java 2
 - net-protocol/all-Java 2
- jdbc.drivers 11
- JNDI
 - context information 26
 - using 22
- JNDI for naming databases 85

L

- LANGUAGE connection property 6, 15
- LANGUAGE_CURSOR 121
- LANGUAGE_CURSOR connection property 15
- Lightweight Directory Access Protocol (LDAP) 23
- LITERAL_PARAMS connection property 16
- localization 32

M

- memory problems in jConnect applications 107
- metadata
 - accessing 47
 - server-side implementation 48
 - USE_METADATA 16
- migrating jConnect applications
 - jConnect applications, migrating 123
- multibyte character sets
 - supported 36
- multibyte character-sets

- converter classes 34
- multithreading
 - making adjustments 95

N

- native-API/partly-Java driver 2
- native-protocol/all-Java driver 2
- net-protocol/all-Java driver 2

O

- Open Server Gateway 22

P

- PACKETSIZE connection property 16
- password 16
- performance, improving 111
 - BigDecimal rescaling 112
 - character-set conversion 113
 - cursors 120
 - tuning for prepared statements in Dynamic SQL 114
- pooling connections 89
- positioned updates and deletes
 - using JDBC 1.x methods 51
 - using JDBC 2.0 methods 52
- PRELOAD_JARS connection property 16
- preloading JAR files 84
- PreparedStatement
 - using with cursors 55
- properties
 - driver 12
- PROTOCOL_CAPTURE connection property 17
- PROXY connection property 17
- PureConverter class 34

R

- Related documents vii
- remote procedure calls (RPCs)

- server-to-server 45
- REMOTEPWD connection property 17
- REPEAT_READ 112
- REPEAT_READ connection property 17
- REQUEST_HA_SESSION 17
- resuming
 - TDS sessions 138
- RMNAME connection property 17
- rows
 - deleting from a cursor result set 54
 - inserting in a cursor result set 55
- rs.getBytes() 68

S

- sample programs 165
- SECONDARY_SERVER_HOSTPORT connection property 18
- SELECT_OPENS_CURSOR connection property 18
- selecting a character-set converter class 35
- sending image data 63
- SERIALIZE_REQUESTS connection property 18
- server-to-server remote procedure calls 45
- SERVICENAME connection property 18
- servlet arguments
 - Debug 136
 - SkipDoneProc 136
 - TdsResponseSize 136
 - TdsSessionIdleTimeout 136
- servlets 127
 - TDS 127
- SESSION_ID connection property 18
- SESSION_TIMEOUT connection property 19
- setRemotePassword() 45
- setting
 - jConnect connection properties 11
 - TDS servlet arguments 136
- setting up
 - jConnect 5
- SkipDoneProc servlet argument 136
- SQL exception and warning messages 141
- SQLINITSTRING connection properties 19
- Statement.cancel() method 10
- stored procedures

- errors 108
- executing 99
 - updating the database from the result set 61
- storing Java objects as column data in a table 76
 - prerequisites 76
 - receiving Java objects from the database 78
 - sending Java objects to a database 77
- STREAM_CACHE_SIZE connection property 19
- Sybase extension changes 125
- SybEventHandler 69
- SybMessageHandler 73
- SYB SOCKET_FACTORY connection property 19
- Syntax conventions ix
- system properties
 - jdbc.drivers 11

T

- TDS 2
 - capturing communication 104
 - installing servlets 136
 - resuming sessions 138
 - servlet system requirements 135
 - servlets 127
 - setting servlet arguments 136
 - tracking sessions 137
 - tunnelling 127
- TdsResponseSize servlet argument 136
- TdsSessionIdleTimeout servlet argument 136
- tracking TDS sessions 137
- troubleshooting 101
- TruncationConverter class 34, 39
- tunnelling
 - TDS 127
- turning off debugging in your application 102
- turning on debugging in your application 102
- TYPE_SCROLL_INSENSITIVE limitations 57

U

- unichar 6
- unichar and univarchar datatypes 33
- updating
 - database from the result set of a stored procedure

Index

61

URL

connection property parameters 21

syntax 20

USE_METADATA connection property 16

user 19

utilities

IsqlApp 163

V

VERSIONSTRING connection property 19

W

Web server gateways 127

wide tables 46

X

XAServer 91