



**PUBLIC**

SDK for SAP Adaptive Server Enterprise 16.0 SP03 PL06

Document Version: 1.0 – 2019-1-15

# SAP jConnect for JDBC Programmers Reference

# Content

<b>1</b>	<b>SAP jConnect for JDBC. . . . .</b>	<b>6</b>
1.1	Java Database Connectivity (JDBC). . . . .	6
<b>2</b>	<b>Programming Information. . . . .</b>	<b>8</b>
2.1	SAP jConnect Version Property. . . . .	8
	SybDriver. setVersion Method. . . . .	9
	JCONNECT_VERSION Connection Property. . . . .	9
2.2	Invoking the SAP jConnect Driver. . . . .	11
2.3	Configuring SAP jConnect for J2EE servers. . . . .	12
2.4	Establish a Connection. . . . .	13
	Connection Properties. . . . .	14
	Connect to SAP ASE. . . . .	37
	Use the sql.ini and Interfaces File Directory Services. . . . .	39
	Connecting to a Server Using JNDI. . . . .	40
2.5	Internationalization and Localization. . . . .	45
	Using SAP jConnect to Pass Unicode Data. . . . .	45
	SAP jConnect Character Set Converters. . . . .	47
2.6	Database Issues. . . . .	53
	Failover Support. . . . .	54
	Server-to-Server Remote Procedure Calls. . . . .	57
	Wide Table Support for Adaptive Server. . . . .	58
	Accessing Database Metadata. . . . .	59
	Use Cursors with Result Sets. . . . .	60
	Support for Batch Updates. . . . .	71
	Updating a Database from a Result Set of a Stored Procedure. . . . .	72
	Datatypes. . . . .	73
2.7	Variable-Length Rows in Data-Only-Locked Tables. . . . .	81
2.8	Large Object (LOB) Support. . . . .	81
2.9	Large Object Locator Support. . . . .	81
2.10	Advanced Features in SAP jConnect. . . . .	82
	BCP Insert. . . . .	83
	Supported SAP ASE Cluster Edition Features. . . . .	84
	Event Notification. . . . .	86
	Error Messages. . . . .	88
	Password Encryption. . . . .	92
	Store Java Objects as Column Data in Table. . . . .	95
	Dynamic Class Loading. . . . .	100

	JDBC 4.0 Specifications Support. . . . .	103
	JDBC 3.0 Specifications Support. . . . .	104
	Support for JDBC 2.0 Optional Package Extensions . . . . .	106
	On-demand Command Encryption for SAP ASE Clients. . . . .	115
	Multiversion Concurrency Control. . . . .	118
2.11	Restrictions and Interpretations of JDBC Standards. . . . .	121
	Unsupported JDBC 4.0 Specification Requirements. . . . .	121
	Use Connection.isClosed and IS_CLOSED_TEST. . . . .	121
	Statement.close with Unprocessed Results. . . . .	122
	Adjustments for Multithreading. . . . .	123
	ResultSet.setCursorName. . . . .	123
	Execute Stored Procedures. . . . .	123
<b>3</b>	<b>Security. . . . .</b>	<b>125</b>
3.1	Restrictions. . . . .	125
3.2	Implement Custom SSL Socket Plug-ins. . . . .	126
	Using Custom Socket with jConnect. . . . .	127
	Create and Configure a Custom Socket. . . . .	127
	SSL Support in SAP jConnect. . . . .	129
	TLS 1.1 and 1.2 Support in SAP jConnect. . . . .	130
3.3	Kerberos. . . . .	130
	Configuring Kerberos for SAP jConnect. . . . .	131
	GSSMANAGER_CLASS Connection Property. . . . .	132
	Kerberos Environment. . . . .	135
	Sample Applications. . . . .	139
	Interoperability. . . . .	141
	Troubleshooting Kerberos. . . . .	142
3.4	Related Documents. . . . .	143
3.5	Securely Automate Database Administration Tasks Using Secure User Store. . . . .	143
	How Does the aseuserstore Utility Work?. . . . .	144
	Enabling Secure User Store for jConnect Client. . . . .	144
	Examples of Creating User Keys for jConnect Clients. . . . .	145
<b>4</b>	<b>Troubleshooting. . . . .</b>	<b>146</b>
4.1	Debugging with SAP jConnect. . . . .	146
	Obtaining an Instance of the Debug Class . . . . .	147
	Turning On Debugging in an Application. . . . .	147
	Turning Off Debugging in an Application. . . . .	148
	Setting the CLASSPATH for Debugging. . . . .	148
	Using the Debugging Methods. . . . .	149
4.2	Dynamic Logging. . . . .	150
	Enabling Logging Dynamically in SAP jConnect. . . . .	151

	Enabling Logging Statically in SAP jConnect. . . . .	152
4.3	Capture TDS Communication. . . . .	152
	PROTOCOL_CAPTURE Connection Property. . . . .	153
	Pause and Resume Methods in Capture Class. . . . .	153
4.4	Resolve Connection Errors. . . . .	154
4.5	Manage Memory in SAP jConnect Applications. . . . .	155
4.6	Resolve Stored Procedure Errors. . . . .	156
	RPC Returns Fewer Output Parameters Than Registered. . . . .	156
	Fetch/State Error . . . . .	156
	Stored Procedure Executed in Unchained Transaction Mode . . . . .	157
4.7	Resolve Custom Socket Implementation Error. . . . .	157
<b>5</b>	<b>Performance and Tuning. . . . .</b>	<b>158</b>
5.1	Improve SAP jConnect Performance. . . . .	158
	BigDecimal Rescaling. . . . .	159
	REPEAT_READ Connection Property. . . . .	159
	SunIoConverter Character-Set Conversion. . . . .	159
5.2	Performance Tuning for Prepared Statements in Dynamic SQL. . . . .	160
	Choose Prepared Statements and Stored Procedures. . . . .	161
	Prepared Statements in Portable Applications . . . . .	162
	Prepared Statements with SAP jConnect Extensions. . . . .	162
	Connection.PrepareStatement. . . . .	164
	DYNAMIC_PREPARE Connection Property. . . . .	164
	SybConnection.PrepareStatement Method . . . . .	165
	ESCAPE_PROCESSING_DEFAULT Connection Property. . . . .	165
	Optimized Batch in SAP jConnect. . . . .	166
5.3	Cursor Performance. . . . .	167
	LANGUAGE_CURSOR Connection Property. . . . .	167
<b>6</b>	<b>Migrating SAP jConnect Applications. . . . .</b>	<b>168</b>
6.1	Migrating Applications to SAP jConnect 16.x. . . . .	168
6.2	Change SAP jConnect Extensions . . . . .	169
	Extension Change Example. . . . .	169
	Method Names. . . . .	170
	Debug Class. . . . .	170
<b>7</b>	<b>Web Server Gateways. . . . .</b>	<b>172</b>
7.1	TDS tunnelling. . . . .	172
7.2	Configure SAP jConnect and Gateways. . . . .	173
	Web Server and SAP ASE on One Host. . . . .	173
	Dedicated JDBC Web Server and SAP ASE on One Host. . . . .	173
	Web Server and SAP ASE on Separate Hosts. . . . .	174

	Connect to Server Through Firewall. . . . .	175
7.3	Usage Requirements. . . . .	175
	Viewing the Index.html File. . . . .	176
	Running Sample Applet. . . . .	176
	Modifying Applet Screen Dimensions. . . . .	177
7.4	TDS-Tunnelling Servlet. . . . .	177
	Reviewing Requirements. . . . .	178
	Installing and Setting Servlet Arguments. . . . .	179
	Invoking the Servlet. . . . .	180
	Tracking Active TDS Sessions. . . . .	180
	Terminating TDS Sessions. . . . .	180
	Resuming a TDS Session. . . . .	181
<b>8</b>	<b>SAP jConnect Sample Programs. . . . .</b>	<b>182</b>
8.1	Running IsqlApp. . . . .	182
<b>9</b>	<b>SAP jConnect Sample Programs and Code. . . . .</b>	<b>185</b>
9.1	Sample Applications . . . . .	185
	Running the Sample Applets. . . . .	186
	Running the Sample Programs with SAP SQL Anywhere. . . . .	186
9.2	Sample Code. . . . .	186
<b>10</b>	<b>SQL Exception and Warning Messages. . . . .</b>	<b>189</b>
<b>11</b>	<b>Glossary. . . . .</b>	<b>214</b>

# 1 SAP jConnect for JDBC

SAP jConnect for JDBC is the SAP high-performance JDBC driver.

SAP jConnect for JDBC is both:

- A native-protocol or all-Java driver, and
- A net-protocol or all-Java driver.

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

- SAP ASE
- SAP SQL Anywhere
- SAP IQ
- SAP Replication Server
- SAP DirectConnect

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

In some instances, the SAP jConnect implementation of JDBC deviates from the JDBC specifications.

## Related Information

[Java Database Connectivity \(JDBC\) \[page 6\]](#)

[Restrictions and Interpretations of JDBC Standards \[page 121\]](#)

## 1.1 Java Database Connectivity (JDBC)

Java Database Connectivity (JDBC), from the Oracle Corporation, 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 and the JDBC Standard Extension API so you can open connections to databases, execute SQL commands, and process results.

## JDBC Interfaces

Interface	Description
<code>java.sql.Driver</code>	Locates the driver for a database URL
<code>java.sql.Connection</code>	Connects to a specific database
<code>java.sql.Statement</code>	Allows users to execute SQL statements.
<code>java.sql.PreparedStatement</code>	Handles SQL statements with parameters
<code>java.sql.CallableStatement</code>	Handles database stored procedure calls
<code>java.sql.ResultSet</code>	Gets the results of SQL statements
<code>java.sql.DatabaseMetaData</code>	Use this interface to access information about the database you have obtained connection to.
<code>java.sql.ResultSetMetaData</code>	Use this interface to retrieve information about ResultSet.
<code>javax.sql.Rowset</code>	Handles JDBC RowSet implementations
<code>javax.sql.DataSource</code>	Handles connection to a data source
<code>javax.sql.ConnectionPoolDataSource</code>	Handles connection pools

Each relational database management system requires a driver to implement these interfaces. 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 reside 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.

For more information about JDBC and its specification, see the [Oracle Technology Network for Java](#) .



## 2 Programming Information

Review the basic components of, and programming requirements for SAP jConnect for JDBC.

Start the SAP jConnect driver, set connection properties, connect to a database server, and review information about using SAP jConnect features. For information about JDBC programming, go to the resource page for Java developers at the [Oracle Technology Network for Java](#) ➔ .

### Related Information

[SAP jConnect Version Property \[page 8\]](#)  
[Invoking the SAP jConnect Driver \[page 11\]](#)  
[Configuring SAP jConnect for J2EE servers \[page 12\]](#)  
[Establish a Connection \[page 13\]](#)  
[Internationalization and Localization \[page 45\]](#)  
[Database Issues \[page 53\]](#)  
[Variable-Length Rows in Data-Only-Locked Tables \[page 81\]](#)  
[Large Object \(LOB\) Support \[page 81\]](#)  
[Large Object Locator Support \[page 81\]](#)  
[Advanced Features in SAP jConnect \[page 82\]](#)  
[Restrictions and Interpretations of JDBC Standards \[page 121\]](#)

### 2.1 SAP jConnect Version Property

The `JCONNECT_VERSION` connection property determines the behavior of the driver and the activated features.

For example, SAP Adaptive Server (SAP ASE) 15.5 supports both SAP jConnect 6.05 and 7.0, however, these two versions process `datetime` and `time` data differently. When connecting to SAP ASE 15.5, SAP jConnect 7.0, which supports microsecond granularity for time data, uses `bigdatetime` or `bigtime` even if the target SAP ASE columns are defined as `datetime` or `time`. SAP jConnect 6.05, however, does not support microsecond granularity and always transfers `datetime` or `time` data when connecting to SAP ASE 15.5.

You can set the SAP jConnect version by using either the `SybDriver.setVersion` method or the `JCONNECT_VERSION` connection property.



## Related Information

[SybDriver.setVersion Method \[page 9\]](#)

[JCONNECT\\_VERSION Connection Property \[page 9\]](#)

### 2.1.1 SybDriver.setVersion Method

The `setVersion` method affects the SAP jConnect default behavior for all connections created by the `SybDriver` object.

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 is made. Changing the version setting during a session does not affect current connections. You can use the `com.sybase.jdbcx.SybDriver.VERSION_LATEST` constant to ensure that you are always requesting the highest version value possible for the SAP jConnect driver you are using. However, by setting the version to `com.sybase.jdbcx.SybDriver.VERSION_LATEST`, you may see changes in behavior if you replace your current SAP jConnect driver with a newer one.

The following code sample shows how to load the SAP jConnect driver and set its version:

```
import java.sql.DriverManager;
import com.sybase.jdbcx.SybDriver;
SybDriver sybDriver = (SybDriver)
    Class.forName("com.sybase.jdbc4.jdbc.SybDriver")
        .newInstance();
sybDriver.setVersion(com.sybase.jdbcx.SybDriver.
    VERSION_7);
DriverManager.registerDriver(sybDriver);
```

### 2.1.2 JCONNECT\_VERSION Connection Property

Use the `JCONNECT_VERSION` connection property to override the `SybDriver` version setting and specify a different version setting for a specific connection.

See the valid `JCONNECT_VERSION` values and the jConnect characteristics associated with these values.

## Features Associated with SAP jConnect Version

JCONNECT_VERSION	Features
7.0	<p>SAP jConnect 7.0 behaves in the same way as SAP jConnect 6.05, except that in 7.0, SAP jConnect requests support for the following:</p> <ul style="list-style-type: none"> <li>• <code>bigdatetime</code> and <code>bigtime</code> SQL datatypes from the server. Versions of SAP Adaptive Server (SAP ASE) earlier than 15.5 ignore this request.</li> <li>• JDBC 4.0.</li> <li>• Valid values of <code>ENABLE_BULK_LOAD</code> are null (default), <code>ARRAYINSERT_WITH_MIXED_STATEMENTS</code>, <code>ARRAYINSERT</code>, <code>BCP</code>, and <code>LOG_BCP</code>.</li> </ul>
6.05	<p>SAP jConnect 6.05 behaves in the same way as SAP jConnect 6.0, except that in 6.05, SAP jConnect requests support for the following:</p> <ul style="list-style-type: none"> <li>• Computed columns, including metadata.</li> <li>• Larger identifiers. With large identifiers, you can use identifiers or object names with lengths of up to 255 bytes. The large identifier applies to most user-defined identifiers, including table name, column name, and index name.</li> </ul>
6.0	<p>SAP jConnect 6.0 behaves in the same way as SAP jConnect 5.x, except that in 6.0, jConnect requests support for the following:</p> <ul style="list-style-type: none"> <li>• <code>date</code> and <code>time</code> SQL datatypes. Versions of SAP ASE earlier than 12.5.1 ignore this request.</li> <li>• <code>unichar</code> and <code>univarchar</code> datatypes from the server. Versions of SAP ASE earlier than 12.5.1 ignore this request.</li> <li>• Wide tables from the server. Versions of SAP ASE earlier than 12.5.1 ignore this request.</li> <li>• Default value of <code>DISABLE_UNICHAR_SENDING</code> is false.</li> </ul>
5.0	<p>SAP jConnect 5.x behaves in the same way as SAP jConnect 4.0.</p>
4.0	<p>SAP jConnect 4.0 behaves in the same way as SAP jConnect 3.0, except that in 4.0, SAP jConnect requests support for the following:</p> <ul style="list-style-type: none"> <li>• The default value of the <code>LANGUAGE</code> connection property is null.</li> <li>• The default behavior of <code>Statement.cancel</code> is to cancel only the <code>Statement</code> object on which it is invoked. This behavior is JDBC-compliant. Use <code>CANCEL_ALL</code> to set the behavior of <code>Statement.cancel</code>.</li> <li>• You can use JDBC 2.0 methods to store and retrieve Java objects as column data.</li> </ul>
3.0	<p>SAP jConnect 3.0 behaves in the same way as SAP jConnect 2.0, except that in 3.0:</p> <ul style="list-style-type: none"> <li>• If the <code>CHARSET</code> connection property does not specify a character set, jConnect uses the default character set of the database.</li> <li>• The default value for <code>CHARSET_CONVERTER</code> is the <code>CheckPureConverter</code> class.</li> </ul>

## JCONNECT\_VERSION

### Features

- |     |   |
|-----|---|
| 2.0 | <ul style="list-style-type: none"><li>• The default value of the <code>LANGUAGE</code> connection property is <code>us_english</code>.</li><li>• If the <code>CHARSET</code> connection property does not specify a character set, the default character set is <code>iso_1</code>.</li><li>• The default value for <code>CHARSET_CONVERTER</code> is the <code>TruncationConverter</code> class, unless the <code>CHARSET</code> connection property specifies a multibyte or 8-bit character set, in which case the default <code>CHARSET_CONVERTER</code> is the <code>CheckPureConverter</code> class.</li><li>• The default behavior of <code>Statement.cancel</code> is to cancel the object it is invoked on and any other <code>Statement</code> objects that have begun to execute and are waiting for results. This behavior is not JDBC-compliant.<br/>Use <code>CANCEL_ALL</code> to set the behavior of <code>Statement.cancel</code>.</li></ul> |
|-----|---|

## Related Information

[JDBC 4.0 Specifications Support \[page 103\]](#)  
[Restrictions and Interpretations of JDBC Standards \[page 121\]](#)  
[SAP jConnect Character Set Converters \[page 47\]](#)  
[Date and Time Datatypes \[page 78\]](#)  
[JDBC 3.0 Specifications Support \[page 104\]](#)  
[Wide Table Support for Adaptive Server \[page 58\]](#)  
[Store Java Objects as Column Data in Table \[page 95\]](#)  
[Using SAP jConnect to Pass Unicode Data \[page 45\]](#)

## 2.2 Invoking the SAP jConnect Driver

Register and invoke SAP jConnect, and add SAP jConnect to the `jdbc.drivers` system property.

### Context

At initialization, the `DriverManager` class attempts to load the drivers listed in `jdbc.drivers`. This is less efficient than calling `Class.forName`. You can list multiple drivers in this property, separated with a colon (:).

#### ❖ Example

This sample code shows how to add a driver to `jdbc.drivers` within a program:

```
Properties sysProps = System.getProperties();  
String drivers = "com.sybase.jdbc4.jdbc.SybDriver";
```

```
String oldDrivers =
sysProps.getProperty("jdbc.drivers");
if (oldDrivers != null)
    drivers += ":" + oldDrivers;
sysProps.put("jdbc.drivers", drivers.toString());
```

### Note

You cannot use `System.getProperties` for Java applets. Use the `Class.forName` method instead.

In Java 6 and JDBC 4, you can use the Java system property `jdbc.drivers` to specify driver classes, for example:

```
java -Djdbc.drivers=com.sybase.jdbc4.jdbc.SybDriver UseDriver
```

You need not use the `UseDriver` program to load the driver explicitly:

```
public class UseDriver
{
    public static void main(String[] args)
    {
        try {
            Connection conn = java.sql.DriverManager.getConnection
                ("jdbc:sybase:Tds:localhost:5000?USER=sa&PASSWORD=secret");
            // more code to use connection ...
        }
        catch (SQLException se){
            System.out.println("ERROR: SQLException "+se);
        }
    }
}
```

## 2.3 Configuring SAP jConnect for J2EE servers

Use the `com.sybase.jdbc4.jdbc.SybConnectionPoolDataSource` class to configure connection pools to the SAP Adaptive Server in application servers (such as EAServer).

### Context

The `com.sybase.jdbc4.jdbc.SybConnectionPoolDataSource` implementation of the `javax.sql.ConnectionPoolDataSource` interface provides `getter` and `setter` methods for every connection property.

#### ❖ Example

You can also configure SAP jConnect programmatically, for example:

```
private DataSource getDataSource ()
{
    SybConnectionPoolDataSource connectionPoolDataSource = new
        SybConnectionPoolDataSource();
```

```

        connectionPoolDataSource.setDatabaseName("pubs2");
        connectionPoolDataSource.setNetworkProtocol("Tds");
        connectionPoolDataSource.setServerName("localhost");
        connectionPoolDataSource.setPortNumber(5000);
        connectionPoolDataSource.setUser("sa");
        connectionPoolDataSource.setPassword(PASSWORD);
        return connectionPoolDataSource;
    }
    private void work () throws SQLException
    {
        Connection conn = null;
        Statement stmt = null;
        DataSource ds = getDataSource();
        try {
            conn = ds.getConnection();
            stmt = conn.createStatement();
            // ...
        }
        finally {
            if (stmt != null) {
                try { stmt.close(); } catch (Exception ex) { /* ignore */ }
            }
            if (conn != null) {
                try { conn.close(); } catch (Exception ex) { /* ignore */ }
            }
        }
    }
}

```

## 2.4 Establish a Connection

Establish a connection to an SAP Adaptive Server or SAP SQL Anywhere database using SAP jConnect.

### Related Information

[Connection Properties \[page 14\]](#)

[Connect to SAP ASE \[page 37\]](#)

[Use the sql.ini and Interfaces File Directory Services \[page 39\]](#)

[Connecting to a Server Using JNDI \[page 40\]](#)

## 2.4.1 Connection Properties

Connection properties specify the information needed to log in to a server and define expected client and server behavior. Connection property names are not case-sensitive.

### Related Information

[Setting Connection Properties \[page 14\]](#)

[Current Connection Settings \[page 15\]](#)

[SAP jConnect Connection Properties \[page 15\]](#)

### 2.4.1.1 Setting Connection Properties

You must set connection properties before you connect to a server.

#### Context

Set the connection properties by either:

- Using the `DriverManager.getConnection` method in your application, or,
- Setting the connection properties when you define the URL.

#### Note

Driver connection properties that you set in the URL do not override any corresponding connection properties set in the application using the `DriverManager.getConnection` method.

#### Example

This sample code uses the `DriverManager.getConnection` method. The sample programs provided with SAP 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);
```

### 2.4.1.2 Current Connection Settings

To view a driver's current connection settings, use `Driver.getDriverPropertyInfo(String url, Properties props)`.

This code returns an array of `DriverPropertyInfo` objects containing:

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

### 2.4.1.3 SAP jConnect Connection Properties

The connection properties for SAP jConnect and their default values.

These properties are not case-sensitive.

You can use the `getClientInfo()` and `setClientInfo()` standard methods to dynamically set the properties indicated as such.



Property	Description
ALTERNATE_SERVER_NAME	<p>Specifies the alternate server name used by the primary and secondary database in a mirrored SAP SQL Anywhere environment. The primary and secondary database use the same alternate server name so that client applications can connect to the current primary server without knowing in advance which of the two servers is the primary server.</p> <p>The JDBC URL syntax is <code>jdbc:sybase:Tds:&lt;hostname&gt;:&lt;port#&gt;/database?connection_property=value;</code>. However, when <code>ALTERNATE_SERVER_NAME</code> is set, SAP jConnect ignores the values of the <code>&lt;hostname&gt;</code> and <code>&lt;port&gt;</code> variables. Instead, SAP jConnect uses the SAP SQL Anywhere UDP discovery protocol to determine the current primary server.</p> <p>For information about database mirroring, see the <i>SAP SQL Anywhere Server - Database Administration Guide</i>.</p> <p>You can also use <code>ALTERNATE_SERVER_NAME</code> with an SAP SQL Anywhere server that is not mirrored. However, you will always get the same host and port values from the singleton server.</p> <p>Default value is null.</p> <p>This property is static.</p>
APPLICATIONNAME	<p>Specifies an application name. This is a user-defined property. You can program the server side to interpret the value provided to this property.</p> <p>Default value is null.</p> <p>This property is static.</p>
BE_AS_JDBC_COMPLIANT_AS_POSSIBLE	<p>Adjusts other properties to ensure that SAP jConnect methods respond in a way that is as compliant as possible with the JDBC 3.0 standard.</p> <p>These properties are affected (and overridden) when this property is set to true:</p> <ul style="list-style-type: none"> <li>• <code>CANCEL_ALL</code> (set to false)</li> <li>• <code>LANGUAGE_CURSOR</code> (set to false)</li> <li>• <code>SELECT_OPENS_CURSOR</code> (set to true)</li> <li>• <code>FAKE_METADATA</code> (set to true)</li> <li>• <code>GET_BY_NAME_USES_COLUMN_LABEL</code> (set to false)</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>

Property	Description
CACHE_COLUMN_METADATA	<p>If you repeatedly use <code>PreparedStatement</code> or <code>CallableStatement</code> objects that perform <code>SELECT</code> queries, setting <code>CACHE_COLUMN_METADATA</code> to true might improve performance. When set to true, the statement remembers the <code>ResultSet</code> metadata information associated with the <code>SELECT</code> query results from the first execution of the statement. On subsequent executions, the metadata is re-used without being reconstructed. This saves CPU time through the use of additional memory.</p> <p>Use the <code>SUPPRESS_ROW_FORMAT</code> connection property when connecting to SAP Adaptive Server (SAP ASE) 15.7 ESD #1 and later.</p> <p>Default value is false.</p> <p>This property is static.</p>
PRE_CACHE_DATATYPE_INFO	<p>If you repeatedly use <code>Statement</code> or its derived interfaces to obtain datatype metadata, setting <code>PRE_CACHE_DATATYPE_INFO</code> to true might improve performance.</p> <p>When <code>PRE_CACHE_DATATYPE_INFO</code> is set to true, information about all user-defined datatypes that serve various <code>ResultSetMetadata</code> APIs, like <code>isCaseSensitive</code> and <code>isSearchable</code>, is cached at connection time. Subsequent access to this information is then available from the cache.</p> <p>When <code>PRE_CACHE_DATATYPE_INFO</code> is false (the default), SAP jConnect does not cache any user-defined datatype information.</p> <div> <p><b>i Note</b></p> <p>Depending on the number of user-defined datatypes that exist in the database to which the connection is being obtained, the time it takes to establish the connection may increase.</p> </div> <p>Default value is false.</p> <p>This property is dynamic.</p>

Property	Description
CANCEL_ALL	<p>Specifies the behavior of the <code>Statement.cancel</code> method:</p> <ul style="list-style-type: none"> <li>• If <code>CANCEL_ALL</code> is false, invoking <code>Statement.cancel</code> cancels only the <code>Statement</code> object on which it is invoked. Thus, if <code>stmtA</code> is a <code>Statement</code> object, <code>stmtA.cancel</code> cancels the execution of the SQL statement contained in <code>stmtA</code> in the database, but no other statements are affected. <code>stmtA</code> is canceled whether it is in cache waiting to execute or has started to execute and is waiting for results.</li> <li>• If <code>CANCEL_ALL</code> is true, invoking <code>Statement.cancel</code> cancels not only the object on which it is invoked, but also any other <code>Statement</code> objects on the same connection that have executed and are waiting for results.</li> </ul> <p>This example sets <code>CANCEL_ALL</code> to false. <code>&lt;props&gt;</code> is a <code>Properties</code> object for specifying connection properties:</p> <pre>props.put("CANCEL_ALL", "false");</pre> <p>To cancel the execution of all <code>Statement</code> objects on a connection, regardless of whether or not they have begun execution on the server, use the extension method <code>SybConnection.cancel</code>.</p> <p>Default values are:</p> <ul style="list-style-type: none"> <li>• true – for <code>JCONNECT_VERSION</code> &lt;= "3"</li> <li>• false – for <code>JCONNECT_VERSION</code> &gt;= "4"</li> </ul> <p>This property is static.</p>

Property	Description
CAPABILITY_TIME	<p>Used only when <code>JCONNECT_VERSION</code> <math>\geq</math> 6. When SAP jConnect is connected to a server that supports the <code>TIME</code> datatype, and all parameters of type <code>java.sql.Time</code> or escape literals <code>{t ...}</code> are processed as <code>TIME</code>. Versions of SAP jConnect earlier than 6.0 treat such parameters as <code>DATETIME</code> and prepend '1970-01-01' to the <code>java.sql.Time</code> parameter. If the underlying datatype is <code>datetime</code> or <code>smalldatetime</code> the date part also gets stored in the database. In SAP jConnect 6.0 or later, when <code>TIME</code> is processed, the server converts time to the underlying datatype and prepends its own base year. This result in incompatibilities between old and new data. If you are using <code>datetime</code> or <code>smalldatetime</code> datatypes for <code>java.sql.Time</code>, then for backward compatibility, leave <code>CAPABILITY_TIME</code> as false. Leaving this property as false forces SAP jConnect to process <code>java.sql.Time</code> parameters or escape literals <code>{t ...}</code> as <code>DATETIME</code> regardless of the server capability of handling <code>TIME</code> datatype. Setting this property to true causes SAP jConnect to process <code>java.sql.Time</code> parameters as <code>TIME</code> datatype when connected to SAP ASE. SAP recommends that you leave this property as false if you are using <code>smalldatetime</code> or <code>datetime</code> columns to store time values.</p> <p>Default value is false.</p> <p>This property is static.</p>
CAPABILITY_WIDETABLE	<p>If you do not require <code>JDBC ResultSetMetaData</code> like Column name as a performance improvement, set this property to false. The result is that less data is exchanged over the network, which improves performance. Unless you are using <code>EAServer</code>, SAP recommends that you use the default setting.</p> <p>Default value is false.</p> <p>This property is static.</p>

Property	Description
CHARSET	<p>Specifies the character set for strings passed to the database. If the <code>CHARSET</code> value is null, SAP jConnect uses the default character set of the server to send <code>string</code> data to the server. If you specify a <code>CHARSET</code>, the database must be able to handle characters in that format. If the database cannot do so, a message is generated indicating that character conversion cannot be properly completed.</p> <p>If you are using SAP jConnect 6.05 or later and the <code>DISABLE_UNICHAR_SENDING</code> is set to false, SAP jConnect detects 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, SAP jConnect sends the character data to the server as <code>unichar</code> data, which allows clients to insert Unicode data into <code>unichar/univarchar</code> columns and parameters.</p> <p>Default value is null.</p> <p>This property is static.</p>
CHARSET_CONVERTER_CLASS	<p>Specifies the character set converter class you want SAP jConnect to use. SAP jConnect uses the version setting from <code>SybDriver.setVersion</code>, or the version passed in with the <code>JCONNECT_VERSION</code> property, to determine the default character-set converter class to use.</p> <p>Default value is version dependent.</p> <p>This property is static.</p>
CLASS_LOADER	<p>A property you set to an <code>DynamicClassLoader</code> object that you create. The <code>DynamicClassLoader</code> loads Java classes that are stored in the database but are not in the <code>CLASSPATH</code> at application start-up time.</p> <p>Default value is null.</p> <p>This property is static.</p>
CONNECTION_FAILOVER	<p>Used with the Java Naming and Directory Interface (JNDI).</p> <p>Default value is true.</p> <p>This property is static.</p>
CRC	<p>When this property is set to true, the update counts that are returned are cumulative counts that include updates directly affected by the statement executed and any triggers invoked as a result of the statement being executed.</p> <p>Default value is false.</p> <p>This property is static.</p>

Property	Description
DATABASE	<p>Use this property to specify the database name for a connection when the connection information is obtained from a <code>SAP interfaces</code> file. The URL of an <code>interfaces</code> file cannot supply the database name.</p> <p>Default value is null.</p> <p>This property is static.</p>
DEFAULT_QUERY_TIMEOUT	<p>When this connection property is set, it is used as the default query timeout for any statements created on this connection.</p> <p>Default value is 0 (no timeout).</p> <p>This property is dynamic.</p>
DELETE_WARNINGS_FROM_EXCEPTION_CHAIN	<p>Values:</p> <ul style="list-style-type: none"> <li>• <code>true</code> – SAP jConnect removes <code>SQLWarning</code> objects from the <code>SQLException</code> chain.</li> <li>• <code>false</code> – SAP jConnect retains the <code>SQLWarning</code> objects in the <code>SQLException</code> chain.</li> </ul> <p>Default value is <code>true</code>.</p> <p>This property is static.</p>
DISABLE_UNICHAR_SENDING	<p>When a client application sends <code>unichar</code> characters to the server (along with non-<code>unichar</code> characters), there is a slight performance hit for any character data sent to the database. This property defaults to <code>false</code> in SAP jConnect 6.05 and later. Clients using older versions of SAP jConnect who want to send <code>unichar</code> data to the database must set this property to <code>false</code>.</p> <p>Default value is version dependent.</p> <p>This property is static.</p>
DISABLE_UNPROCESSED_PARAM_WARNINGS	<p>Disables warnings. During results processing for a stored procedure, SAP jConnect often reads return values other than row data. If you do not process the return value, SAP jConnect raises a warning. To disable these warnings (which might improve performance), set this property to <code>true</code>.</p> <p>Default value is <code>false</code>.</p> <p>This property is static.</p>
DYNAMIC_PREPARE	<p>Specifies whether to retain or remove <code>SQLWarning</code> from the Determines whether dynamic SQL prepared statements are precompiled in the database.</p> <p>Default value is <code>true</code>.</p> <p>This property is dynamic.</p>

Property	Description
EARLY_BATCH_READ_THRESHOLD	<p>Specifies the number of rows after which a reader thread should be started to drain out the server responses for a batch.</p> <p>Set this value to -1 if the early read is never required.</p> <p>Default value is -1.</p> <p>This property is static.</p>
ENABLE_BULK_LOAD	<p>Specifies whether to use bulk load to insert rows to the database.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• null – disables bulk load.</li> <li>• ARRAYINSERT_WITH_MIXED_STATEMENTS – enables bulk load with row-level logging and allows your application to execute other statements during the bulk-load operation.</li> <li>• ARRAYINSERT – enables bulk load with row-level logging, but your application cannot execute other statements during the bulk-load operation.</li> <li>• BCP – enables bulk load with page-level logging; your application cannot execute other statements during the bulk-load operation.</li> <li>• LOG_BCP – same as BCP except the complete transaction is dumped for possible full recovery.</li> </ul> <p>Default value is null.</p> <p>This property is dynamic.</p>
ENABLE_LOB_LOCATORS	<p>Specifies whether SAP jConnect should create a client-side materialized LOB or server-side LOB locator.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• false – SAP jConnect uses client-side materialized LOBs. That is, the entire LOB data is processed and cached on the client side.</li> <li>• true – works only when autocommit is set to false, otherwise internally, the value changes to false. When set to true, server locators are used instead of storing the LOB data on client side.</li> </ul> <p>Default value is false.</p> <p>This property is dynamic.</p>
ENABLE_SERVER_PACKETSIZE	<p>Specifies if the connection packet size is set to the value suggested by the server. If true, the driver does not use <code>PACKETSIZE</code> connection property, and the server is free to use any value between 512 and the maximum packet size. If false, the <code>PACKETSIZE</code> connection property is used.</p> <p>Default value is true.</p> <p>This property is static.</p>



Property	Description
ENCRYPT_PASSWORD	<p>Allows a secure login. When this property is true, both login and remote site passwords are encrypted before being sent to the server. Passwords are no longer sent in clear text.</p> <p>ENCRYPT_PASSWORD has precedence over RETRY_WITH_NO_ENCRYPTION.</p> <p>Default value is false.</p> <p>This property is static.</p>
ESCAPE_PROCESSING_DEFAULT	<p>Circumvents processing of JDBC function escapes in SQL statements. By default, SAP 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 avoid this processing, which might provide a slight performance benefit.</p> <p>Additionally, ESCAPE_PROCESSING_DEFAULT helps with back-end servers such as SAP® IQ that use curly braces as part of the SQL syntax.</p> <p>Default value is true.</p> <p>This property is static.</p>
EXECUTE_BATCH_PAST_ERRORS	<p>Specifies whether SAP jConnect allows a batch update operation to ignore nonfatal errors encountered while executing individual statements and to complete the batch update, or aborts the batch update operation.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• true – allows a batch update operation to ignore nonfatal errors encountered and to complete the batch update.</li> <li>• false – aborts a batch update when a nonfatal error is encountered.</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>
EXPIRESTRING	<p>Contains the license expiration date. Expiration is set to Never except for evaluation copies of SAP jConnect.</p> <p>Default value is never.</p> <p>This property is static and read-only.</p>

Property	Description
FAKE_METADATA	<p>Returns fake metadata. When you call the <code>ResultSetMetaData</code> methods <code>getCatalogName</code>, <code>getSchemaName</code>, and <code>getTableName</code> and this property is true, the call returns empty strings ("") because the server does not supply useful metadata.</p> <p>When this property is false, calling these methods throws a "Not Implemented" <code>SQLException</code>.</p> <p>If you have enabled wide tables and are using an SAP ASE 12.5 or later, this property setting is ignored because the server supplies useful metadata.</p> <p>Default value is false.</p> <p>This property is static.</p>
GET_BY_NAME_USES_COLUMN_LABEL	<p>Provides backward compatibility with versions of SAP jConnect earlier than 6.0.</p> <p>With SAP ASE version 12.5 and later, SAP jConnect has access to more metadata than was previously available. Prior to version 12.5, <code>column name</code> and <code>column alias</code> meant the same thing. SAP jConnect can now differentiate between the two when used with a 12.5 or later SAP ASE with wide tables enabled.</p> <p>To preserve backward compatibility, set this property to true. If you want calls to <code>getBytes</code>, <code>getInt</code>, <code>get*</code> (<code>String columnName</code>) to look at the actual name for the column, set this property to false.</p> <p>Default value is true.</p> <p>This property is static.</p>
GET_COLUMN_LABEL_FOR_NAME	<p>Maintains backward compatibility with SAP jConnect 5.5 or earlier, where a call to <code>ResultSetMetaData.getColumnLabel</code> returns the column label rather than the column name.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>true</code> – <code>ResultSetMetaData.getColumnLabel</code> returns column label.</li> <li>• <code>false</code> – <code>ResultSetMetaData.getColumnLabel</code> returns column name.</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>

Property	Description
GSSMANAGER_CLASS	<p>Specifies a third-party implementation of the <code>org.ietf.jgss.GSSManager</code> class.</p> <p>You can set this property to a string or a <code>GSSManager</code> object.</p> <p>If you set the property to a string, the value should be the fully qualified class name of the third-party <code>GSSManager</code> implementation. If you set the property to an object, the object must extend the <code>org.ietf.jgss.GSSManager</code> class.</p> <p>Default value is null.</p> <p>This property is static.</p>
HOMOGENEOUS_BATCH	<p>Invokes the SAP ASE optimized batching protocol to speed up batch operations for <code>PreparedStatement</code> objects.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>true</code> – optimized batching protocol is used.</li> <li>• <code>false</code> – unoptimized batching protocol is used even if SAP jConnect is connected to an SAP ASE that supports new optimized batching protocol.</li> </ul> <p>Default value is <code>true</code>.</p> <p>This property is dynamic.</p>
HOSTNAME	<p>Identifies the name of the current host.</p> <p>Default value is none. The max length is 30 characters and, if exceeded, it is truncated to 30.</p> <p>This property is static.</p>
HOSTPROC	<p>Identifies the application process on the host machine.</p> <p>Default value is none.</p> <p>This property is static.</p>
IGNORE_DONE_IN_PROC	<p>Intermediate update results (as in stored procedures) are not returned; only the final result set is.</p> <p>Default value is <code>false</code>.</p> <p>This property is static.</p>

Property	Description
IGNORE_WARNINGS	<p>Specifies whether or not to check for and generate warning messages. This property checks only for warnings regarding the loss of precision when storing timestamp values into SAP ASE <code>date</code> and <code>time</code> datatypes, which have lower precision than the Java timestamp.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>true</code> – SAP jConnect does not check for and generate warning messages, thus improving performance.</li> <li>• <code>false</code> – the default value, which directs SAP jConnect to check and generate warning messages.</li> </ul> <p>Before setting <code>IGNORE_WARNINGS</code> to <code>true</code>, thoroughly test the impact of such a configuration on your application.</p> <p>Default value is <code>false</code>.</p> <p>This property is static.</p>
IMPLICIT_CURSOR_FETCH_SIZE	<p>Use this property with the <code>SELECT_OPENS_CURSOR</code> property to force SAP jConnect to open a read-only cursor on every <code>select</code> query that is sent to the database. The cursor has a fetch size of the value set in this property, unless overridden by the <code>Statement.setFetchSize</code> method.</p> <p>Default value is 0.</p> <p>This property is static.</p>
INTERNAL_QUERY_TIMEOUT	<p>Use this property to set the query timeout used by statements internally created and executed by SAP jConnect. This may prevent application failures if internal commands do not complete in a reasonable time.</p> <p>Default value is 0(no timeout).</p> <p>This property is dynamic.</p>
IS_CLOSED_TEST	<p>Allows you to specify what query, if any, is sent to the database when <code>Connection.isClosed</code> is called.</p> <p>Default value is null.</p> <p>This property is static.</p>
J2EE_TCK_COMPLIANT	<p>When this property is <code>true</code>, the SAP jConnect driver enables behavior that is compliant with the J2EE 1.4 technology compatibility kit (TCK) test suite, which causes some loss of performance. Therefore, SAP recommends using the default value of <code>false</code>.</p> <p>Default value is <code>false</code>.</p> <p>This property is static.</p>

Property	Description
JAVA_CHARSET_MAPPING	<p>Specifies a user-defined character set mapping that supersedes the default SAP ASE character set mapping.</p> <p>Default value is none.</p> <p>This property is static.</p>
JCE_PROVIDER_CLASS	<p>Specifies the Java Cryptography Extension (JCE) provider used in RSA encryption algorithms.</p> <p>Default value is bundled JCE provider.</p> <p>This property is static.</p>
JCONNECT_VERSION	<p>Sets version-specific characteristics.</p> <p>Default value is 7.</p> <p>This property is static.</p>
LANGUAGE	<p>Specifies the language in which messages from SAP jConnect and the server appear. The setting must match a language in <code>syslanguages</code> because server messages are localized according to the language setting in your local environment. The languages supported are Chinese, US English, French, German, Japanese, Korean, Polish, Portuguese, and Spanish.</p> <p>Default value is version dependent.</p> <p>This property is static.</p>
LANGUAGE_CURSOR	<p>Determines that SAP jConnect uses language cursors instead of protocol cursors.</p> <p>Default value is false.</p> <p>This property is static.</p>
LITERAL_PARAMS	<p>When true, 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 false, parameter markers are left in the SQL statement and the parameter values are sent to the server separately.</p> <p>Default value is false.</p> <p>This property is static.</p>
NEWPASSWORD	<p>Specifies the new password used in password expiration handling.</p> <p>Default value is null.</p> <p>This property is static.</p>

Property	Description
OPTIMIZE_FOR_PERFORMANCE	<p>Specifies whether or not to enable SAP jConnect performance-enhancing properties. This property controls only the IGNORE_WARNINGS property.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• true – SAP jConnect runs in enhanced performance mode.</li> <li>• false – the default value, which means that SAP jConnect runs in normal mode.</li> </ul> <p>Before setting OPTIMIZE_FOR_PERFORMANCE to true, thoroughly test the impact of such a configuration on your application.</p> <p>Default value is false.</p> <p>This property is static.</p>
OPTIMIZE_STRING_CONVERSIONS	<p>Specifies whether or not to enable string conversion optimization.</p> <p>This optimization behavior might improve SAP jConnect performance when a client uses character datatypes in SQL prepared statements.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• 0 – string conversion optimization is not enabled.</li> <li>• 1 – enables string conversion optimization when SAP jConnect uses UTF8 or server default character set.</li> <li>• 2 – enables string conversion optimization for all cases.</li> </ul> <p>Default value is 0.</p> <p>This property is static.</p>
PACKETSIZE	<p>Identifies the network packet size. If you are using SAP ASE 15.0 or later, SAP recommends that you do not set this property, and allow SAP jConnect and SAP ASE to select the network packet size that is appropriate for your environment.</p> <p>Default value is 512.</p> <p>This property is static.</p>
PASSWORD	<p>Identifies the login password. <code>String</code>.</p> <p>Set automatically if using the <code>getConnection(String, String, String)</code> method, or explicitly if using <code>getConnection(String, Props)</code>.</p> <p>Default value is none.</p> <p>This property is static.</p>

Property	Description
PRELOAD_JARS	<p>Contains a comma-separated list of <code>.jar</code> file names that are associated with the <code>CLASS_LOADER</code> that you specify. These <code>.jar</code> files are loaded at connect time, and are available for use by any other connection using the same SAP jConnect driver.</p> <p>Default value is null.</p> <p>This property is static.</p>
PROMPT_FOR_NEWPASSWORD	<p>Specifies whether to perform a transparent password change or to prompt for the new password.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>true</code> – prompts you to manually set the new password.</li> <li>• <code>false</code> – SAP jConnect checks the value of <code>NEWPASSWORD</code> and, if it is not null, uses this value to replace the expired password.</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>
PROTOCOL_CAPTURE	<p>Specifies a file for capturing TDS communication between an application and an SAP ASE.</p> <p>Default value is null.</p> <p>This property is dynamic.</p>
PROXY	<p>Specifies a gateway address. For the HTTP protocol, the URL is <code>http://host:port</code>.</p> <p>To use the HTTPS protocol that supports encryption, the URL is <code>https://host:port/servlet_alias</code>.</p> <p>Default value is none.</p> <p>This property is static.</p>
QUERY_TIMEOUT_CANCELS_ALL	<p>Forces SAP jConnect to cancel all statements on a connection when a read timeout is encountered. This behavior can be used when a client has calls <code>execute()</code> and the timeout occurs because of a deadlock (for example, trying to read from a table that is currently being updated in another transaction).</p> <p>Default value is false.</p> <p>This property is dynamic.</p>



Property	Description
RELEASE_LOCKS_ON_CURSOR_CLOSE	<p>Specifies whether SAP ASE releases shared read-only cursor locks at isolation levels 2 and 3 when a cursor is closed:</p> <ul style="list-style-type: none"> <li>• false – does not enable shared cursor locks release on close.</li> <li>• true – enables shared cursor locks release on close.</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>
REMOTEPWD	<p>Contains remote server passwords for access through server-to-server remote procedure calls.</p> <p>Default value is none.</p> <p>This property is static.</p>
REPEAT_READ	<p>Determines whether the driver keeps copies of columns and output parameters so that columns can be read out of order or repeatedly.</p> <p>Default value is true.</p> <p>This property is static.</p>
REQUEST_HA_SESSION	<p>Indicates whether the connecting client wants to begin a high availability (HA) failover session.</p> <p>You cannot reset the property once a connection has been made. For additional flexibility for requesting failover sessions, code the client application to set REQUEST_HA_SESSION at runtime.</p> <p>Setting this property to true causes SAP jConnect to attempt a failover login. If you do not set this connection property, a failover session does not start, even if the server is configured for failover.</p> <p>Default value is false.</p> <p>This property is static.</p>
REQUEST_KERBEROS_SESSION	<p>Determines whether SAP jConnect uses Kerberos for authentication. If you set this property to true, you must also enter a value for the SERVICE_PRINCIPAL_NAME property.</p> <p>You may also want to provide a value for the GSSMANAGER_CLASS property.</p> <p>Default value is false.</p> <p>This property is static.</p>

Property	Description
RETRY_WITH_NO_ENCRYPTION	<p>Allows server to retry logging in using clear text passwords.</p> <p>When both <code>ENCRYPT_PASSWORD</code> and <code>RETRY_WITH_NO_ENCRYPTION</code> properties are set to true, SAP jConnect first logs in using the encrypted password. If login fails, SAP jConnect logs in using the clear text password.</p> <p>Default value is false.</p> <p>This property is static.</p>
RMNAME	<p>Sets the Resource Manager name when using distributed transactions (XA). This property overrides a Resource Manager name that may be set in an LDAP server entry.</p> <p>Default value is null.</p> <p>This property is static.</p>
SECURE_STORE_KEY	<p>Specifies the key that retrieves the user's credentials from the secure store. It is a static property that should be specified before establishing connection with the SAP ASE server.</p>
SECONDARY_SERVER_HOSTPORT	<p>Sets the host name 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 true.</p> <p>Default value is null.</p> <p>This property is static.</p>
SELECT_OPENS_CURSOR	<p>Determines whether calls to <code>Statement.executeQuery</code> automatically generate a cursor when the query contains a <code>FOR UPDATE</code> clause.</p> <p>If you have previously called <code>Statement.setFetchSize</code> or <code>Statement.setCursorName</code> on the same statement, a setting of true for <code>SELECT_OPENS_CURSOR</code> has no effect.</p> <p>You may experience some performance degradation when <code>SELECT_OPENS_CURSOR</code> is set to true.</p> <p>Default value is false.</p> <p>This property is static.</p>

Property	Description
SEND_BATCHPARAMS_IMMEDIATE	<p>Specifies whether SAP jConnect sends the parameters for the current row immediately after invoking <code>PreparedStatement.addBatch()</code> or only after invoking <code>PreparedStatement.executeBatch()</code>.</p> <ul style="list-style-type: none"> <li>• <code>true</code> – SAP jConnect sends the parameters for the current row immediately after invoking <code>PreparedStatement.addBatch()</code>. This minimizes usage of client memory and gives the server more time to process the batch parameters.</li> <li>• <code>false</code> – SAP jConnect sends the batch parameters only after invoking <code>PreparedStatement.executeBatch()</code>.</li> </ul> <p>Default value is <code>false</code>.</p> <p>This property is dynamic.</p>
SERIALIZE_REQUESTS	<p>Determines whether SAP jConnect waits for responses from the server before sending additional requests.</p> <p>Default value is <code>false</code>.</p> <p>This property is static.</p>
SERVER_INITIATED_TRANSACTIONS	<p>Allows the server to control transactions. By default the property is set to <code>true</code> and SAP jConnect lets the server start and control transactions by using the SAP Transact-SQL command <code>set chained on</code>. If set to <code>false</code>, the transactions are started and controlled by SAP jConnect by using the Transact-SQL command <code>begin tran</code>. SAP recommends that you allow the server to control the transactions.</p> <p>Default value is <code>true</code>.</p> <p>This property is static.</p>
SERVICENAME	<p>Indicates the name of a back-end database server that a DirectConnect gateway serves. Also used to indicate which database to use upon connecting to SAP SQL Anywhere.</p> <p>Default value is <code>none</code>.</p> <p>This property is static.</p>
SERVERTYPE	<p>When connected to an SAP® Open Switch™, set this property to <code>OSW</code>, which allows SAP jConnect to send certain instructions to SAP Open Switch that allows SAP Open Switch to remember initial connection settings for example, isolation level, textsize, quoted identifier and autocommit when SAP Open Switch redirects a connection to a different server instance.</p> <p>Default value is <code>none</code>.</p> <p>This property is static.</p>

Property	Description
SERVICE_PRINCIPAL_NAME	<p>Used when establishing a Kerberos connection to SAP ASE. The value of this property should correspond both to the server entry in your key distribution center (KDC) and to the server name under which your database is running.</p> <p>The value of the <code>SERVICE_PRINCIPAL_NAME</code> property is ignored if the <code>REQUEST_KERBEROS_SESSION</code> property is set to false.</p> <p>Default value is null.</p> <p>This property is static.</p>
SESSION_ID	<p>A TDS session ID. When this property is set, SAP jConnect assumes that an application is trying to resume communication on an existing TDS session held open by the TDS-tunnelling gateway. SAP jConnect skips the login negotiations and forwards all requests from the application to the specified session ID.</p> <p>Default value is null.</p> <p>This property is static.</p>
SESSION_TIMEOUT	<p>Specifies the amount of time, in seconds, that an HTTP-tunnelled session (created using the SAP jConnect TDS-tunnelling servlet) remains alive while idle. After the specified time, the connection is automatically closed.</p> <p>Default value is null.</p> <p>This property is static.</p>
SETMAXROWS_AFFECTS_SELECT_ONLY	<p>Specifies whether <code>setMaxRows</code> limits only the rows returned by <code>select</code> statements to be consistent with the JDBC specification.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>true</code> – <code>Statement.setMaxRows(int max)</code> limits only the number of rows returned as a result of the <code>select</code> statements.</li> <li>• <code>false</code> – <code>Statement.setMaxRows(int max)</code> limits the number of rows returned as a result of the <code>select</code>, <code>insert</code>, <code>update</code>, and <code>delete</code> statements.</li> </ul> <p><code>SETMAXROWS_AFFECTS_SELECT_ONLY</code> is ignored when connected to SAP ASE 15.5 or earlier.</p> <p>Default value is true.</p> <p>This property is static.</p>
SQLINITSTRING	<p>Defines 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.</p> <p>Default value is null.</p> <p>This property is static.</p>

Property	Description
STREAM_CACHE_SIZE	<p>Specifies the maximum size used to cache statement response streams.</p> <p>Default value is null (unlimited cache size).</p> <p>This property is dynamic.</p>
STRIP_BLANKS	<p>Forces the server to remove the preceding and trailing blanks in a string value before storing it in the table.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• false – string values sent by the client are stored as is.</li> <li>• true – preceding and trailing blanks in a string value are removed before storing it in the table.</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>
SUPPRESS_CONTROL_TOKEN	<p>Suppresses control tokens.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• false – control tokens are sent.</li> <li>• true – control tokens are suppressed.</li> </ul> <p>Default value is false.</p> <p>This property is static.</p>
SUPPRESS_PARAM_FORMAT	<p>When executing dynamic SQL prepared statements, SAP jConnect client can use the <code>SUPPRESS_PARAM_FORMAT</code> connection string property to suppress parameter format metadata. The client sends less parameter metadata where possible for better performance.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• false – parameter format metadata is not suppressed in select, insert, and update operations.</li> <li>• true – the default value; parameter format metadata is suppressed where possible.</li> </ul> <p>Default value is true.</p> <p>This property is static.</p>

Property	Description
SUPPRESS_ROW_FORMAT	<p>In SAP jConnect, client can use the <code>SUPPRESS_ROW_FORMAT</code> connection string property to force SAP ASE to send <code>TDS_ROWFM1</code> or <code>TDS_ROWFM2</code> data only when the row format changes for a dynamic SQL prepared statement. SAP ASE can send less data to the client if possible, resulting in better performance.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>false</code> – <code>TDS_ROWFM1</code> or <code>TDS_ROWFM2</code> data is sent, even if the row format has not changed.</li> <li>• <code>true</code> – the default; forces the server to send <code>TDS_ROWFM1</code> or <code>TDS_ROWFM2</code> only when the row format has changed.</li> </ul> <p>Default value is <code>true</code>.</p> <p>This property is static.</p>
SUPPRESS_ROW_FORMAT2	<p>Specifies that SAP ASE is to send data using the <code>TDS_ROWFM1</code> byte sequence where possible instead of the <code>TDS_ROWFM2</code> byte sequence.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>• <code>false</code> – the default value; <code>TDS_ROWFM2</code> is not suppressed.</li> <li>• <code>true</code> – forces the server to send data in <code>TDS_ROWFM1</code> where possible.</li> </ul> <p>When connected to SAP ASE 15.7 ESD #1 or later, use the <code>SUPPRESS_ROW_FORMAT</code> connection property instead.</p> <p>Default value is <code>false</code>.</p> <p>This property is static.</p>
SYB SOCKET_FACTORY	<p>Enables SAP jConnect to use your custom socket implementation.</p> <p>Set <code>SYB SOCKET_FACTORY</code> either to:</p> <ul style="list-style-type: none"> <li>• The name of a class that implements <code>com.sybase.jdbcx.SybSocketFactory</code>; or</li> <li>• <code>DEFAULT</code>, which instantiates a new <code>java.net.Socket ( )</code></li> </ul> <p>Use this property to make an SSL connection to your database.</p> <p>Default value is <code>null</code>.</p> <p>This property is static.</p>

Property	Description
TEXTSIZE	<p>Allows you to set the text size. By default, SAP ASE and SAP SQL Anywhere allow 32,627 bytes to be read from an image or text column. If you have the SAP jConnect Meta Data tables installed, SAP jConnect changes that value to 2GB. However, setting this value when connected to SAP Open Switch allows the connection to remember the setting when SAP Open Switch redirects a connection to a different server instance.</p> <p>Default value is 2GB.</p> <p>This property is static.</p>
USE_METADATA	<p>Creates and initializes a <code>DatabaseMetaData</code> object when you establish a connection. The <code>DatabaseMetaData</code> object is necessary to connect to a specified database.</p> <p>SAP 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 SAP jConnect. See <i>Installing Stored Procedures</i> in the <i>jConnect for JDBC Installation Guide</i>.</p> <p>Default value is true.</p> <p>This property is static.</p>
USER	<p>Specifies the login ID.</p> <p>Set automatically if using the <code>getConnection(String, String, String)</code> method, or explicitly if using <code>getConnection(String, Props)</code>.</p> <p>Default value is none.</p> <p>This property is static.</p>
VERSIONSTRING	<p>Provides read-only version information for the JDBC driver.</p> <p>Default value is SAP jConnect driver version.</p> <p>This property is static.</p>

## Related Information

[DYNAMIC\\_PREPARE Connection Property \[page 164\]](#)

[Password Encryption \[page 92\]](#)

[Security \[page 125\]](#)

[Optimized Batch in SAP jConnect \[page 166\]](#)

[Cursor Performance \[page 167\]](#)  
[Failover Support \[page 54\]](#)  
[Wide Table Support for Adaptive Server \[page 58\]](#)  
[CONNECTION\\_FAILOVER Property \[page 43\]](#)  
[Large Object Locator Support \[page 81\]](#)  
[Use Connection.isClosed and IS\\_CLOSED\\_TEST \[page 121\]](#)  
[Supersede Default Character Set Mapping \[page 52\]](#)  
[JCONNECT\\_VERSION Connection Property \[page 9\]](#)  
[Release Locks at Cursor Close \[page 67\]](#)  
[TDS tunnelling \[page 172\]](#)  
[Using DynamicClassLoader \[page 100\]](#)  
[Using SAP jConnect to Pass Unicode Data \[page 45\]](#)  
[Selecting a Character Set Converter \[page 48\]](#)  
[Preloading .jar Files \[page 102\]](#)

## 2.4.2 Connect to SAP ASE

In Java application, define a URL using the SAP jConnect driver to connect to an SAP Adaptive Server Enterprise (SAP ASE) server.

### Context

The basic format of the URL is:

```
jdbc:sybase:Tds:<host>:<port>
```

where:

- `<jdbc:sybase>` identifies the driver.
- `<Tds>` is the SAP communication protocol for SAP ASE.
- `<host:port>` is the SAP ASE server host name and listening port. See `$SYBASE/interfaces (UNIX)` or `%SYBASE%\ini\sql.ini (Windows)` for the entry that your database or SAP 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>
```

#### i Note

Connect to a specific database using SAP SQL Anywhere or DirectConnect. Use the `SERVICENAME` connection property to specify the database name instead of `"/database"`.



### ❖ Example

The following code creates a connection to an SAP ASE 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);
```

## Related Information

[URL Connection Property Parameters \[page 38\]](#)

### 2.4.2.1 URL Connection Property Parameters

Specify the values for the SAP jConnect driver connection properties when you define a URL.

#### i Note

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

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>
```

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"
```

## 2.4.3 Use the sql.ini and Interfaces File Directory Services

Use the `sql.ini` file for Windows and the `interfaces` file for UNIX to provide server information for SAP jConnect for JDBC.

### Context

By using the `sql.ini` or `interfaces` file, enterprises can centralize the information about the services available in the enterprise networks including SAP ASE information.

Use the connection string to identify the `sql.ini` or `interfaces` file. On jConnect for JDBC, you can connect to only a single Directory Services URL (DSURL).

### Related Information

[Connection String for Single DSURL for SAP jConnect \[page 39\]](#)

[Format of sql.ini and Interfaces Files for SSL \[page 40\]](#)

### 2.4.3.1 Connection String for Single DSURL for SAP jConnect

When connecting to a DSURL, you must specify the path to the `sql.ini` or `interfaces` file and the server name.

If you do not to set the path, SAP jConnect returns an error.

This specifies the path to the `sql.ini` file:

```
String url = "jdbc:sybase:jndi:file:///D:/syb1252/ini/mysql.ini?myaseISO1"
```

where:

- server name = myaseISO1
- `sql.ini` file path = `file:///D:/syb1252/ini/sql.ini?`

This specifies the path to the `interfaces` file:

```
String url = "jdbc:sybase:jndi:file:///work/sybase/interfaces?myase"
```

where:

- server name = myase
- `interfaces` file path = `file:///work/sybase/interfaces`

## 2.4.3.2 Format of sql.ini and Interfaces Files for SSL

Review the format of `sql.ini` and `interfaces` files for SSL.

Format for `sql.ini` file for SSL:

```
[SYBSRV2]
master=nlwmsck,mango1,4100,ssl
query=nlwmsck,mango1,4100,ssl
query=nlwmsck,mango1,5000,ssl
```

The format for the `interfaces` file is:

```
sybsrv2
master tcp ether mango1 5000 ssl
query tcp ether mango1 4100 ssl
query tcp ether mango1 5000 ssl
```

### Note

SAP jConnect supports multiple query entries under the same server name in the `sql.ini` or `interfaces` file. SAP jConnect attempts to connect to values for `host` or `port` from the query entry in the sequence, as in the `sql.ini` or `interfaces` file. If SAP jConnect finds a SSL in a query entry, it requires the application to be coded to handle SSL connections by specifying an application specific socket factory, or the connection may fail.

## 2.4.4 Connecting to a Server Using JNDI

In SAP jConnect, use the Java Naming and Directory Interface (JNDI) to provide connection information.

### Context

SAP jConnect provides:

- 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 SAP jConnect `CONNECTION_FAILOVER` property for handling unsuccessful connection attempts. When `CONNECTION_FAILOVER` is true, SAP jConnect attempts to connect to a sequence of host/port server addresses in the JNDI name space until one succeeds.

Using SAP jConnect with JNDI, 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.

## Related Information

[Connection URL for Using JNDI \[page 41\]](#)

[Required Directory Service Information \[page 41\]](#)

[CONNECTION\\_FAILOVER Property \[page 43\]](#)

[Providing JNDI Context Information \[page 44\]](#)

[Connection URL for Using JNDI \[page 41\]](#)

[Required Directory Service Information \[page 41\]](#)

[CONNECTION\\_FAILOVER Property \[page 43\]](#)

[Providing JNDI Context Information \[page 44\]](#)

### 2.4.4.1 Connection URL for Using JNDI

To specify that SAP jConnect use Java Naming and Directory Interface (JNDI) to obtain connection information, place “jndi” as the URL protocol after “sybase”.

For example:

```
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.

### 2.4.4.2 Required Directory Service Information

Review the required directory service information when using JNDI with SAP jConnect.

JNDI must return this 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

Stores this information 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.

#### i Note

You can use the alias name to reference the attribute instead of the OID.

## Directory Service Information for JNDI

Attribute Description	Alias	OID (object_id)
Interfaces entry replacement in LDAP directory services	sybaseServer	1.3.6.1.4.1.8974.1.1
Collection point for sybaseServer LDAP attributes	sybaseServer	1.3.6.1.4.1.8974.2
Version	sybaseVersion	1.3.6.1.4.1.8974.2.1
Server name	sybaseServer	1.3.6.1.4.1.8974.2.2
Service	sybaseService	1.3.6.1.4.1.8974.2.3
Status	sybaseStatus	1.3.6.1.4.1.8974.2.4
(Required)Address	sybaseAddress	1.3.6.1.4.1.8974.2.5
Security mechanism	sybaseSecurity	1.3.6.1.4.1.8974.2.6
Retry count	sybaseRetryCount	1.3.6.1.4.1.8974.2.7
Loop delay	sybaseRetryDelay	1.3.6.1.4.1.8974.2.8
(Required)jConnect connection protocol	sybaseJconnectProtocol	1.3.6.1.4.1.8974.2.9
(Required)jConnect connection property	sybaseJconnectProperty	1.3.6.1.4.1.8974.2.10
(Required)Database name	sybaseDatabasename	1.3.6.1.4.1.8974.2.11
High availability failover servername	sybaseHAservername	1.3.6.1.4.1.8974.2.15
ResourceManager name	sybaseResourceManagerName	1.3.6.1.4.1.8974.2.16
ResourceManager type	sybaseResourceManagerType	1.3.6.1.4.1.8974.2.17
JDBCDataSource interface	sybaseJdbcDataSource- Interface	1.3.6.1.4.1.8974.2.18
ServerType	sybaseServerType	1.3.6.1.4.1.8974.2.19

These examples show connection information entered for the database server “SYBASE11” under an LDAP directory service. You can use either the OID or the alias.

**Example 1** uses the attribute OID:

```
dn: servername=SYBASE11,o=MyCompany,c=US
   servername:SYBASE11
1.3.6.1.4.1.8974.2.5:TCP#1#giotto 1266
1.3.6.1.4.1.8974.2.5:TCP#1#giotto 1337
1.3.6.1.4.1.8974.2.5:TCP#1#standby1 4444
1.3.6.1.4.1.8974.2.10:REPEAT_READ=false&
   PACKETSIZE=1024
1.3.6.1.4.1.8974.2.10: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** uses the attribute alias, which is not case sensitive:

```
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 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”.

### 2.4.4.3 CONNECTION\_FAILOVER Property

CONNECTION\_FAILOVER is a boolean-valued connection property you can use when SAP jConnect uses JNDI to get connection information.

If CONNECTION\_FAILOVER is true (the default), SAP 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, SAP 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, if a database server is associated with these 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, SAP jConnect tries to connect to the host “giotto” at port 1266. If this fails, SAP jConnect tries port 1337 on “giotto.” If this fails, SAP jConnect tries to connect to host “standby1” through port 4444.

If CONNECTION\_FAILOVER is false, SAP jConnect attempts to connect to an initial host and port number. If the attempt fails, SAP jConnect throws a SQL exception and does not try again.

## 2.4.4.4 Providing JNDI Context Information

Be familiar with the JNDI specification to use SAP jConnect with JNDI.

### Context

See the [JNDI specification from Oracle Technology Network](#) .

In particular, make sure that required initialization properties are set in `javax.naming.directory.DirContext` when JNDI and SAP jConnect are used together. Set these properties either at the system level or at runtime.

The properties are:

- `Context.INITIAL_CONTEXT_FACTORY` – 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` – 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`”.

#### ❖ Example

This example shows how to set context properties at runtime and how to get a connection using JNDI and LDAP. The `INITIAL_CONTEXT_FACTORY` context property is set to invoke the Oracle implementation of an LDAP service provider. The `Context.PROVIDER_URL` property is set to the URL of an LDAP directory service located on the host “`ldap_server1`” at port 389.

```
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_server1:389");

/* 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_server1:389" +
    "/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, SAP jConnect passes them to JNDI to be used in initializing a server, as in this SAP jConnect code:

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

SAP jConnect then obtains the connection information it needs from JNDI by invoking `DirContext.getAttributes`, as in this example, where `<ctx>` is a `DirContext` object:

```
javax.naming.directory.Attributes attrs =  
    ctx.getAttributes("ldap://ldap_server1:389/servername=" +  
        "Sybase11", SYBASE_SERVER_ATTRIBUTES);
```

`SYBASE_SERVER_ATTRIBUTES` is an array of strings defined within SAP jConnect. The array values are the OIDs for the required directory information listed in [Required Directory Service Information \[page 41\]](#).

## 2.5 Internationalization and Localization

Review the internationalization and localization issues relevant to SAP jConnect.

### Related Information

[Using SAP jConnect to Pass Unicode Data \[page 45\]](#)

[SAP jConnect Character Set Converters \[page 47\]](#)

### 2.5.1 Using SAP jConnect to Pass Unicode Data

In SAP Adaptive Server (SAP ASE) version 12.5 and later, database clients can take advantage of the `unichar` and `univarchar` datatypes

#### Context

The two datatypes allow for the efficient storage and retrieval of Unicode data, allowing users to designate database table columns to store Unicode data, regardless of the default character set of the server.

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.

### **i Note**

In SAP ASE version 12.5 through 12.5.0.3, the server was required to use a default character set of utf-8 to use the Unicode datatypes. However, in SAP ASE 12.5.1 and later, database users can use `unichar` and `univarchar` without having to consider the default character set of the server.

You can use the `unichar` and `univarchar` datatypes anywhere that you can use `char` and `varchar` character datatypes, without having to make syntax changes.

- `unichar` – use `<n>` to specify the number of Unicode characters (the amount of storage allocated is 2 bytes per character).
- `univarchar` – use `<n>` to specify the maximum length in characters for the variable-length datatypes.

When the server accepts `unichar` and `univarchar` data, SAP jConnect behaves as follows:

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

For example, if a client attempts to send a Unicode Japanese character to an SAP ASE 12.5.1 that has `iso_1` as the default character set, SAP jConnect detects that the Japanese character cannot be converted to an `iso_1` character. SAP jConnect then sends the string as Unicode data.

There is a performance penalty when a client sends `unichar`/`univarchar` data to a server, because SAP 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.

If you are using a SAP jConnect version that is earlier than 6.05 and you want to use the `unichar` and `univarchar` datatypes, you must:

## **Procedure**

1. Set the `JCONNECT_VERSION = 6` or later.
2. Set the `DISABLE_UNICHAR_SENDING` connection property to false.

For more information on support for `unichar` and `univarchar` datatypes, see *SAP ASE Enterprise Manuals*.

## **Related Information**

[ENABLE\\_UTF8\\_DECODE\\_WARNING Connection Property \[page 47\]](#)

[JCONNECT\\_VERSION Connection Property \[page 9\]](#)

### 2.5.1.1 **ENABLE\_UTF8\_DECODE\_WARNING Connection Property**

The `ENABLE_UTF8_DECODE_WARNING` property applies only when the server charset is UTF8, jConnect uses the default `CharSetConverter`, and an application tries to fetch data.

When the property is set to `TRUE`, jConnect displays a SQL warning for characters that cannot be mapped to UTF characters, and as a result, data gets truncated from the position where the UTF8 characters that cannot be mapped, are found.

By default, the value of the `ENABLE_UTF8_DECODE_WARNING` connection property is set to `FALSE`, and jConnect throws a SQL exception if UTF8 characters that cannot be mapped, are found in the data.

## 2.5.2 **SAP jConnect Character Set Converters**

There are two character set conversion classes. The conversion class that SAP jConnect uses is based on the `JCONNECT_VERSION`, `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. The `TruncationConverter` class is the default converter when `JCONNECT_VERSION` is set to 2.  
Using the `TruncationConverter` class, SAP jConnect 16.0 handles character sets in the same manner as SAP jConnect version 2.2. The `TruncationConverter` class is the default converter when the `JCONNECT_VERSION = 2`.
- The `PureConverter` class is a pure Java, multibyte character-set converter. SAP jConnect uses this class if the `JCONNECT_VERSION = 4` or later. SAP jConnect also uses this converter when `JCONNECT_VERSION = 2` if it detects a character set specified in the `CHARSET` connection property that is incompatible with the `TruncationConverter` class.  
Although it enables multibyte character-set conversions, the `PureConverter` class may negatively impact SAP jConnect driver performance.

## **Related Information**

[Selecting a Character Set Converter \[page 48\]](#)

[Setting the CHARSET Connection Property \[page 48\]](#)

[Improving Character Set Conversion Performance \[page 49\]](#)

[Supported Character Sets \[page 49\]](#)

[European Currency Symbol Support \[page 53\]](#)

## 2.5.2.1 Selecting a Character Set Converter

SAP jConnect uses the `JCONNECT_VERSION` to determine the default character-set converter class to use.

### Context

For `JCONNECT_VERSION` = 2.0 or 3.0, the default is `TruncationConverter`. For `JCONNECT_VERSION` = 4.0 or later, the default is `PureConverter`.

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

#### ❖ Example

For example, if you set `JCONNECT_VERSION` = 4.0 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.jdbc4.charset.TruncationConverter")
```

## 2.5.2.2 Setting the CHARSET Connection Property

Specify the character set to use in your application by setting the `CHARSET` driver property.

### Context

If you do not set the `CHARSET` property:

- For `JCONNECT_VERSION` = 2.0, SAP jConnect uses `iso_1` as the default character set.
- For `JCONNECT_VERSION` = 3.0 through 6.05, SAP jConnect uses the default character set of the database, and adjusts automatically to perform any necessary conversions on the client side.
- For SAP jConnect versions starting with 6.05, if SAP jConnect cannot successfully convert the user data to the negotiated charset, it sends unconverted Unicode characters to the server if the server supports the Unicode characters, otherwise, it throws an exception.

You can also use the `-J <charset>` command line option for the `IsqlApp` application to specify a character set.

### ❖ Example

To determine which character sets are installed on your SAP Adaptive Server Enterprise (SAP ASE), issue this SQL query on your server:

```
select name from syscharsets
go
```

For the `PureConverter` class, if the designated `CHARSET` does not work with the client Java Virtual Machine (JVM), the connection fails with a `SQLException`, indicating that you must set `CHARSET` to a character set that is supported by both, SAP ASE, 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. Therefore, if your application must process non-ASCII data (for instance, any Asian languages), do not use `TruncationConverter`, as this causes data corruption.

## 2.5.2.3 Improving Character Set Conversion Performance

If you use multibyte character sets and need to improve driver performance, you can use the `SunIoConverter` class provided with the SAP jConnect samples.

### Context

In addition, you can use `TruncationConverter` to improve performance if your application deals with only 7-bit ASCII data.

### Related Information

[SunIoConverter Character-Set Conversion \[page 159\]](#)

## 2.5.2.4 Supported Character Sets

Character sets supported by SAP jConnect, and the corresponding JDK byte converter for each supported character set.

Although SAP jConnect supports UCS-2, currently no SAP databases or SAP Open Servers support UCS-2.

SAP Adaptive Server Enterprise (SAP ASE) versions 12.5 and later support a version of Unicode known as UTF-16 encoding.

## Supported SAP jConnect Character Sets

SybCharset Name	JDK Byte Converter
ascii_7	ASCII
big5	Big5
big5hk (for JDK 1.3 and above)	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
cp949	Cp949
cp950	Cp950
cp1250	Cp1250
cp1251	Cp1251
cp1252	Cp1252
cp1253	Cp1253
cp1254	Cp1254
cp1255	Cp1255

SybCharset Name	JDK Byte Converter
cp1256	Cp1256
cp1257	Cp1257
cp1258	Cp1258
deckanji	EUC_JP
eucgb	EUC_CN
eucjis	EUC_JP
eucksc	EUC_KR
gb18030	GB18030
ibm420	Cp420
ibm918	Cp918
iso_1	ISO8859_1
iso88592	ISO8859_2
iso88595	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

SybCharset Name	JDK Byte Converter
ucs2	Unicode
utf8	UTF8

## Related Information

[Unsupported Character Sets \[page 52\]](#)

[Supersede Default Character Set Mapping \[page 52\]](#)

### 2.5.2.4.1 Unsupported Character Sets

Some character sets are not supported in SAP jConnect because no JDK byte converters are analogous to them.

- cp1047
- euccns
- greek8
- roman8
- roman9
- 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.

### 2.5.2.4.2 Supersede Default Character Set Mapping

Use the `JAVA_CHARSET_MAPPING` connection property to supersede the default SAP Adaptive Server Enterprise (SAP ASE) character set mapping.

**Example** maps the server character set cp949 to ms949:

```
props.put("CHARSET", "cp949"); /* Server character set */
props.put("JAVA_CHARSET_MAPPING", "ms949"); /* Java character set
mapping */
```

Most of the SAP ASE character sets have the same name as the Java character sets that they are mapped to. See [Supported Character Sets \[page 49\]](#) for those character sets that are mapped to a Java character set with a different name.

## 2.5.2.5 European Currency Symbol Support

SAP jConnect supports the use of the European currency symbol, or “euro,” and its conversion to and from UCS-2 Unicode.

The euro is included in these SAP jConnect character sets: cp1250, cp1251, cp1252, cp1253, cp1254, cp1255, cp1256, cp1257, cp1258, cp874, iso885915, and utf8.

To use the euro symbol:

- Use the `PureConverter` or `CheckPureConverter` class, pure Java, multibyte character-set converter.
- Verify that the new character sets are installed on the server.
- Select the appropriate character set on the client.

### Related Information

[SAP jConnect Character Set Converters \[page 47\]](#)

[Setting the CHARSET Connection Property \[page 48\]](#)

## 2.6 Database Issues

Review the database issues relevant to SAP jConnect.

### Related Information

[Failover Support \[page 54\]](#)

[Server-to-Server Remote Procedure Calls \[page 57\]](#)

[Wide Table Support for Adaptive Server \[page 58\]](#)

[Accessing Database Metadata \[page 59\]](#)

[Use Cursors with Result Sets \[page 60\]](#)

[Support for Batch Updates \[page 71\]](#)

[Updating a Database from a Result Set of a Stored Procedure \[page 72\]](#)

[Datatypes \[page 73\]](#)

[Support for Batch Updates \[page 71\]](#)

[Datatypes \[page 73\]](#)

[Failover Support \[page 54\]](#)

[Server-to-Server Remote Procedure Calls \[page 57\]](#)

[Wide Table Support for Adaptive Server \[page 58\]](#)

[Use Cursors with Result Sets \[page 60\]](#)

[Transact-SQL Queries with COMPUTE Clause \[page 70\]](#)



[Variable-Length Rows in Data-Only-Locked Tables \[page 81\]](#)

[Large Object \(LOB\) Support \[page 81\]](#)

[Large Object Locator Support \[page 81\]](#)

[Accessing Database Metadata \[page 59\]](#)

[Updating a Database from a Result Set of a Stored Procedure \[page 72\]](#)

## 2.6.1 Failover Support

SAP jConnect supports the SAP Adaptive Server Enterprise (SAP ASE) failover.

SAP ASE Failover allows you to configure two SAP ASE servers as companions.

### **i** Note

SAP ASE Failover in a high availability system is a different feature than connection failover. SAP strongly recommends that you read this section very carefully if you want to use both.

If the primary companion fails, the devices, databases, and connections for that server can be taken over by the secondary companion. You can configure a high availability system either asymmetrically or symmetrically.

- An asymmetric configuration includes two SAP ASE servers that are physically located on different machines but are connected so that if one of the servers is brought down, the other assumes its workload. The secondary SAP ASE acts as a “hot standby” and does not perform any work until failover occurs.
- A symmetric configuration also includes two SAP ASE servers running on separate machines. However, if failover occurs, either SAP ASE can act as a primary or secondary companion for the other SAP ASE. In this configuration, each SAP ASE 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 SAP jConnect and connect a client application to an SAP ASE 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.

See *Using Failover in High Availability Systems in the Adaptive Server Documentation* for more detailed information.

When using SAP jConnect as part of your failover strategy:

- Have two SAP ASE servers configured for failover.
- Only changes that were committed to the database before failover are retained when the client fails over.
- Set the `REQUEST_HA_SESSION` jConnect connection property to true.
- SAP jConnect event notification does not work when failover occurs.
- Close all statements when they are no longer used. SAP jConnect stores information on statements to enable failover. Unclosed statements result in memory leaks.

## Related Information

[Implementing Failover in SAP jConnect \[page 55\]](#)

## 2.6.1.1 Implementing Failover in SAP jConnect

Implement failover support in SAP jConnect.

### Procedure

1. Set:
  - `REQUEST_HA_SESSION` to true.
  - `SECONDARY_SERVER_HOSTPORT` to the host name and port number where your secondary server is listening.
2. Use JNDI to connect to the server. 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 has 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) – assumes that the search base (typically provided by the `java.naming.provider.url` attribute), combined with the value of this attribute, 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) – assumes that the value of the HA attribute 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 – 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
```

Use the `REQUEST_HA_SESSION` connection property to indicate that the connecting client wants to begin a failover session with an SAP ASE server that is configured for failover. Setting this property to true instructs SAP jConnect to attempt a failover login. If you do not set this connection property, a failover session does 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.

This example shows connection information entered for the database server SYBASE1 under an LDAP directory service, where "tahiti" is the primary server, and "moorea" is the secondary companion server:

```
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
```

### 3. Request a connection using JNDI and LDAP:

- Use the directory of the LDAP 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_server1:389" +
"/servername=Sybasell,o=MyCompany,c=US", props);
```

, or

- Specify a searchbase:

```
props.put(Context.PROVIDER_URL,
"ldap://ldap_server1:389/ o=MyCompany, c=US");
```

```
Connection con=DriverManager.getConnection
```

```
("jdbc:sybase:jndi:servername=Sybase11", props);
```

Failover process allows:

**Logging in to the primary server** if an SAP ASE server is not configured for failover or cannot grant a failover session, the client cannot log in.

```
'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 secondary server** when failover occurs, the SQL exception JZOF2 is thrown:

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

The client automatically reconnects to the secondary database using JNDI and lets you:

- Identity 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.
- Restart your application with a procedure or return to the last completed transaction or activity.

**Failing back to primary server** the system administrator determines the timing of failback, by issuing `sp_failback` on the secondary server. The client fails back from the secondary server to the primary server.

After failback, the client can expect the same behavior and results on the primary server during failover to the secondary server.

## Related Information

[Connection Properties \[page 14\]](#)

[Connecting to a Server Using JNDI \[page 40\]](#)

## 2.6.2 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, SAP jConnect uses the password of the current connection 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.

SAP jConnect includes a property that enables you to set a universal remote password or different passwords on several servers.

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 must import the `SybDriver` class, then call the method:

```
import com.sybase.jdbcx.SybDriver;
SybDriver sybDriver = (SybDriver)
    Class.forName("com.sybase.jdbc4.jdbc.SybDriver").newInstance();
sybDriver.setRemotePassword
    (<serverName>, password, connectionProps);
```

#### **i Note**

To set different remote passwords for various servers, repeat the preceding call 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 is 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, SAP jConnect no longer sets the default universal password.

## **2.6.3 Wide Table Support for Adaptive Server**

SAP Adaptive Server Enterprise (SAP ASE) 15.7 ESD #1 offers limits and parameters that are larger than previous versions of the database server.

For example:

- Tables can contain 1024 columns.
- `varchar` and `varbinary` columns can contain more than 255 bytes of data.
- You can send and retrieve up to 2048 parameters when invoking stored procedures or as parameters to `PreparedStatement`.
- When connected to SAP ASE 15.7 ESD #1 and later, you can send and retrieve up to 32767 parameters to `PreparedStatement`.

To ensure that SAP jConnect requests wide table support from the database, the default setting of `JCONNECT_VERSION` must be 6.0 or later.

### i Note

SAP jConnect continues to work with an SAP ASE version 12.5 and later if you set `JCONNECT_VERSION` to earlier than 6.0. 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 `JCONNECT_VERSION` to 6.0 or later when you access data from an SAP ASE server that does not support wide tables. In this case, the server simply ignores your request for wide table support.

In addition to the larger number of columns and parameters, wide table support provides extended result set metadata. For example, in versions of SAP jConnect earlier than 6.0, the `ResultSetMetaData` methods `getCatalogName`, `getSchemaName`, and `getTableName` all returned `NotImplementedSQLExceptions` 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.

## 2.6.4 Accessing Database Metadata

To support `DatabaseMetaData` methods, SAP Adaptive Server Enterprise (SAP ASE) provides a set of stored procedures that SAP jConnect can call to obtain metadata about a database.

### Context

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

- `sql_server.sql` – installs stored procedures on SAP ASE databases earlier than version 12.0.
- `sql_server12.sql` – installs stored procedures on SAP ASE database version 12.0.x.
- `sql_server12.5.sql` – installs stored procedures on SAP ASE database version 12.5.x.
- `sql_server15.0.sql` – installs stored procedures for SAP ASE 15.0 through 15.5.
- `sql_server15.7.sql` – installs stored procedures for SAP ASE 15.7 or 15.7 ESD #1.
- `sql_server15.7.0.2.sql` – installs stored procedures for SAP ASE 15.7 ESD #2 or later.
- `sql_server16.0.sql` – installs stored procedures for SAP ASE 16.0.
- `sql_asa.sql` – installs stored procedures on the SAP SQL Anywhere database version 9.x.
- `sql_asa10.sql` – installs stored procedures on the SAP SQL Anywhere database version 10.x.
- `sql_asa11.sql` – installs stored procedures on the SAP SQL Anywhere database version 11.x.
- `sql_asa12.sql` – installs stored procedures on the SAP SQL Anywhere database version 12.x.
- `sql_asa16.sql` – installs stored procedures on the SAP SQL Anywhere database version 16.x.

### i Note

The most recent versions of these scripts are compatible with all versions of SAP jConnect.

See the *SAP jConnect for JDBC Installation Guide* and *SAP jConnect for JDBC Release Bulletin* for information about installing stored procedures.

### ❖ Example

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 obtain metadata from temporary tables in a database.

### i Note

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

## 2.6.5 Use Cursors with Result Sets

SAP jConnect 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, `ResultSet`s 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 Javadoc.

When requested, SAP jConnect opens server-side scrollable cursors when the server is SAP ASE 15.0 or later.

`java.sql.ResultSet` Options Available in SAP jConnect

Concurrency	Type		
	TYPE_FORWARD_ONLY	TYPE_SCROLL_INSENSITIVE	TYPE_SCROLL_SENSITIVE
<code>CONCUR_READ_ONLY</code>	Supported	Supported	Not available
<code>CONCUR_UPDATABLE</code>	Supported	Not available	Not available

## Related Information

[Cursors \[page 61\]](#)

[JDBC 1.x Methods for Positioned Updates and Deletes \[page 63\]](#)

[JDBC 2.0 Methods for Positioned Updates and Deletes \[page 64\]](#)

[Release Locks at Cursor Close \[page 67\]](#)

[Select for Update Support \[page 67\]](#)

[Cursor with PreparedStatement Object \[page 67\]](#)

[TYPE\\_SCROLL\\_INSENSITIVE Result Sets in SAP jConnect \[page 68\]](#)

[Transact-SQL Queries with COMPUTE Clause \[page 70\]](#)

[JDBC 2.0 Methods for Positioned Updates and Deletes \[page 64\]](#)

[Cursor with PreparedStatement Object \[page 67\]](#)

[TYPE\\_SCROLL\\_INSENSITIVE Result Sets in SAP jConnect \[page 68\]](#)

[JDBC 1.x Methods for Positioned Updates and Deletes \[page 63\]](#)

[Creating and Using a Cursor \[page 62\]](#)

## 2.6.5.1 Cursors

Methods for creating a cursor using SAP jConnect.

- `SybStatement.setCursorName` – assigns explicitly the cursor a name.

The signature for `SybStatement.setCursorName` is:

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

- `SybStatement.setFetchSize` – creates a cursor and specifies 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 SAP jConnect driver names the cursor. To get the name of the cursor, use `ResultSet.getCursorName`.

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

```
Statement createStatement(int resultSetType, int resultSetConcurrency) throws SQLException
```

If you request an unsupported `ResultSet`, a SQL warning is chained to the connection. When the returned `Statement` is executed, you receive the kind of `ResultSet` that is most like the one you requested. See the *JDBC Specification* for more details on the behavior of this method.

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

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

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

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



## Related Information

[Creating and Using a Cursor \[page 62\]](#)

### 2.6.5.1.1 Creating and Using a Cursor

Use the `Statement.setCursorName` or `SybStatement.setFetchSize` method to create a cursor .

#### Procedure

1. Create a 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.

This example uses each of the two methods for creating cursors and returning a result set. It also uses `ResultSet.getCursorName` 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 = 'Pinole'");
while(rs2.next())
{
    ...
}

// Get the name of the cursor created through the
// setFetchSize() method.
String cursor_name = rs2.getCursorName();
...
// For jConnect 6.0, 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();
...

```

## 2.6.5.2 JDBC 1.x Methods for Positioned Updates and Deletes

Review the methods to use JDBC 1.x.

This 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.

```

// 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

Use the Statement object `<stmt2>` to perform a positioned deletion

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

### 2.6.5.3 JDBC 2.0 Methods for Positioned Updates and Deletes

The JDBC 2.0 methods to update the columns in the current cursor row and the database from the current cursor row in a result set.

#### Related Information

[Updating Columns in Result Sets \[page 64\]](#)

[Methods for Updating a Database from a Result Set \[page 65\]](#)

[Deleting Rows from Result Sets \[page 66\]](#)

[Inserting Rows into Result Sets \[page 66\]](#)

#### 2.6.5.3.1 Updating Columns in Result Sets

JDBC 2.0 specifies a number of methods for updating column values from a result set in memory, on the client.

#### Context

You can then use the updated values 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 SAP 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;
```

## 2.6.5.3.2 Methods for Updating a Database from a Result Set

JDBC 2.0 specifies 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 raise exceptions. For `insertRow`, all table columns that require non-null entries must be specified. Methods provided on `DatabaseMetaData` dictate when these changes are visible.

### ❖ Example

#### Example

This example creates a single `Statement` object that returns a cursor result set. For each row in the result set, column values are updated in memory and the database is updated with the new column values for the row.

```
// 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(ResultSet.TYPE_SCROLL_INSENSITIVE,  
        ResultSet.CONCUR_UPDATABLE);  
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 6.0
    Statement stmt = conn.createStatement
        (ResultSet.TYPE_FORWARD_ONLY, ResultSet.CONCUR_UPDATABLE);
    // In jConnect 6.0, 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...");
    }

```

### 2.6.5.3.3 Deleting Rows from Result Sets

Delete a row from a cursor result set.

#### Context

To delete a row, use the `SybCursorResultSet.deleteRow`:

```

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

```

### 2.6.5.3.4 Inserting Rows into Result Sets

Insert a row using the JDBC 2.0 API.

#### Context

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();
```

## 2.6.5.4 Release Locks at Cursor Close

SAP ASE 15.7 extends the `declare cursor` syntax to include the `release_locks_on_close` option, which releases shared cursor locks at isolation levels 2 and 3 when a cursor is closed.

SAP jConnect accordingly supports the release-lock-on-close semantics.

To use SAP jConnect connection, set the `RELEASE_LOCKS_ON_CURSOR_CLOSE` connection property to true. The default value is false.

This setting takes effect only when connected to a server that supports `release_locks_on_close`.

For information about `release_locks_on_close`, see the *SAP ASE Enterprise Reference Manual:Commands*.

## 2.6.5.5 Select for Update Support

SAP ASE 15.7 and later supports `select for update`, which can lock rows for subsequent updates within the same transaction, and supports exclusive locks for cursors that can be updated.

See *Queries: Selecting Data from a Table* in the *SAP ASE Enterprise Transact-SQL Users Guide*.

This functionality is automatically available to clients when the `for update` clause is added to a `select` statement and to any cursors that can be updated, and are opened within the clients.

## 2.6.5.6 Cursor with PreparedStatement Object

You can use `PreparedStatement` multiple times with the same or different values for its input parameters.

If you use a cursor with a `PreparedStatement` object, you must 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`).

This 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

```

## 2.6.5.7 TYPE\_SCROLL\_INSENSITIVE Result Sets in SAP jConnect

SAP jConnect supports TYPE\_SCROLL\_INSENSITIVE result sets.

SAP jConnect uses the Tabular Data Stream (TDS)— its proprietary protocol—to communicate with SAP database servers. SAP Adaptive Server Enterprise (SAP ASE) 15.0 and later supports TDS scrollable cursors. For servers that do not support TDS scrollable cursors, SAP 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 memory. Because this may cause a performance strain, SAP recommends that you use TYPE\_SCROLL\_INSENSITIVE result sets only with SAP ASE 15.0 or when the result set is reasonably small.

### i Note

When you use TYPE\_SCROLL\_INSENSITIVE `ResultSet`s in SAP jConnect, and the server does not support TDS scrollable cursors, you can call the `isLast` method only after the last row of the `ResultSet` has been read. Calling `isLast` before the last row is reached throws an `UnsupportedOperationException`.

SAP 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 insensitive 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 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);
```





Class Diagram Showing Relationship Between Sample Classes and JDBC API

See the JDBC 2.0 API at [Oracle Technology Network for Java](https://docs.oracle.com/javase/7.0/api/java/sql/package-summary.html) for more details.

## 2.6.5.8 Transact-SQL Queries with COMPUTE Clause


SAP jConnect for JDBC supports Transact-SQL queries that include a `COMPUTE` clause.

A `COMPUTE` clause allows you to display detail and summary results in one `select` statement. The summary row appears following the detail rows of a specific group. For example:

```
select type, price, advance
  from titles
 order by type
 compute sum(price), sum(advance) by type
```

type	price	advance
UNDECIDED	NULL	NULL
Compute Result:		
NULL		NULL
type	price	advance
business	2.99	10,125.00
business	11.95	5,000.00
business	19.99	5,000.00
business	19.99	5,000.00
Compute Result:		
54.92		25,125.00
...		
...		
(24 rows affected)		

When SAP jConnect executes a `select` statement that includes a `COMPUTE` clause, SAP jConnect returns multiple result sets to the client. The number of result sets depends on the number of unique groupings available. Each group contains one result set for the detail rows and one result set for the summary. The client must process all result sets to fully process the rows returned; if it does not, only the detail rows of the first group of data are included in the first result set returned.

For more information about the `COMPUTE` clause, see the *SAP ASE Enterprise Transact-SQL Users Guide*. For more information about processing multiple result sets, see the JDBC API documentation on the [Oracle Technology Network for Java](#)  Web site.

## 2.6.6 Support for Batch Updates

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

Any statement added to a batch must return only an update count and cannot return a `ResultSet`.

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

SAP jConnect also supports dynamic `PreparedStatement`s in batch.

### Related Information

[Implementation Notes \[page 71\]](#)


#### 2.6.6.1 Implementation Notes

SAP jConnect implements batch updates as specified in the JDBC 2.0 API.

Exceptions are:

- The `EXECUTE_BATCH_PAST_ERRORS` connection property controls how failures are handled in batch execution.  
By default, `EXECUTE_BATCH_PAST_ERRORS` is false and SAP jConnect stops processing after the first failure. `BatchUpdateException.getUpdateCounts` returns an `int[]` array with 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. "N" represents the total statements in the batch.  
When `EXECUTE_BATCH_PAST_ERRORS` is true, SAP jConnect continues processing in the presence of nonfatal failures. `BatchUpdateException.getUpdateCounts` returns an `int[]` array with length of `N`, where "N" represents the total statements in the batch. Examine the individual update counts to determine execution status of each statement.
- To call stored procedures in batch (unchained) mode, you must create the stored procedure in unchained mode.

- If SAP Adaptive Server Enterprise (SAP ASE) encounters a fatal error during batch execution, `BatchUpdateException.getUpdateCounts` returns only an `int[]` length of zero. The entire transaction is rolled back if a fatal error is encountered, resulting in zero successful rows.
- Batch updates in databases that do not support batch updates: SAP 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.

For details on batch updates, see the [JDBC API documentation](#) .

## Related Information

[Stored Procedure Executed in Unchained Transaction Mode \[page 157\]](#)

## 2.6.7 Updating a Database from a Result Set of a Stored Procedure

SAP jConnect includes `update` and `delete` methods that allow you to get a cursor on the result set returned by a stored procedure.

### Context

You can then use the position of the cursor 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.

### ❖ Example

To get a cursor on the result set returned by a stored procedure, use either `SybCallableStatement.setCursorName` or `SybCallableStatement.setFetchSize` before you execute the callable statement that contains the procedure. This 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}",
        ResultSet.TYPE_FORWARD_ONLY, ResultSet.CONCUR_UPDATABLE);

// 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);
}
```

## 2.6.8 Datatypes

Review the use of numeric, image, text, date, time, and char data.

### Related Information

[Numeric Datatype \[page 73\]](#)

[Image Datatype \[page 74\]](#)

[Text Datatype \[page 77\]](#)

[Date and Time Datatypes \[page 78\]](#)

[Char, Varchar, Text, and GetByte Datatypes \[page 79\]](#)

[Other Supported Datatypes \[page 79\]](#)

### 2.6.8.1 Numeric Datatype

The `SybPreparedStatement` extension supports the way SAP Adaptive Server Enterprise (SAP ASE) 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 SAP 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.

The `SybPreparedStatement` extension used with the following method gives SAP jConnect applications more control over `setBigDecimal`:

```
public void setBigDecimal (int parameterIndex, BigDecimal X, int scale,
```

```
int precision) throws SQLException
```

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

## 2.6.8.2 Image Datatype

SAP jConnect has a `TextPointer` class with `sendData` methods for updating an image column in an SAP Adaptive Server Enterprise (SAP ASE) or SAP SQL Anywhere database.

In versions of SAP jConnect earlier than 4.0, you had to send image data using the `setBinaryStream` method in `java.sql.PreparedStatement`. In version 5.0 and later, the `TextPointer.sendData` methods use `java.io.InputStream` and greatly improve performance when you send image data to an SAP ASE database.

### ⚠ Caution

Using the `TextPointer` class with `sendData()` method may affect the application as `TextPointer` is not a standard JDBC form.

SAP recommends you use `PreparedStatement.setBinaryStream(int paramIndex, InputStream image)` or utilize the LOB locator support, both standard JDBC forms to send image data. However, `setBinaryStream()` may consume much more memory on procedure cache than the `TextPointer` class when large image data is handled.

Until a replacement for the `TextPointer` class is implemented, SAP will continue supporting it.

To obtain instances of the `TextPointer` class, use either of these methods in `SybResultSet`:

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

## Related Information

[Public Methods in TextPointer Class \[page 74\]](#)

### 2.6.8.2.1 Public Methods in TextPointer Class

Review the public methods in `TextPointer` class in SAP jConnect.

The `com.sybase.jdbcx` 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

```

where:

- `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.
- `<log>` is a parameter for each method that specifies whether image data is to be fully logged in the database transaction log. If the `log` parameter is true, the entire binary image is written into the transaction log. If the `log` parameter is false, the update is logged, but the image itself is not included in the log.

## Related Information

[TextPointer Object \[page 75\]](#)

[Executing the Update with TextPointer.sendData \[page 76\]](#)

[Updating an Image Column with TextPointer.sendData \[page 77\]](#)

### 2.6.8.2.1.1 TextPointer Object

The 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, select the column into a `ResultSet` object, then use `SybResultSet.getTextPtr`, which 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. To shortcut this process, use the

set `textsize` command to minimize the amount of data returned in a packet. This 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");
```

## 2.6.8.2.1.2 Executing the Update with `TextPointer.sendData`

Use the `TextPointer` object to update the `pic` column with image data in the file `Anne_Ringer.gif`.

### Context

Sample code:

```
/*
 *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 SAP jConnect installation directory for more information.

### 2.6.8.2.1.3 Updating an Image Column with `TextPointer.sendData`

Update a column with image data using `TextPointer.sendData`.

#### Procedure

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

#### Results

In this 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.

## 2.6.8.3 Text Datatype

In SAP jConnect 3.0 and earlier versions, a `TextPointer` class is used with `sendData` methods for updating a text column in an SAP Adaptive Server Enterprise (SAP ASE) server or SAP SQL 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 SAP ASE server or SAP SQL Anywhere, 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)
```

## 2.6.8.4 Date and Time Datatypes

SAP jConnect for JDBC supports the SAP Adaptive Server Enterprise (SAP ASE) `datetime`, `smalldatetime`, `bigdatetime`, `bigtime`, `date`, and `time` datatypes:

- `datetime` can hold dates between January 1, 1753 and December 31, 9999 that are accurate to 1/300 second on platforms that support this level of granularity.
- `smalldatetime` can hold dates from January 1, 1900 to June 6, 2079, with accuracy to the minute.
- `bigdatetime` indicates the number of microseconds that have passed since January 1, 0000 0:00:00.000000. The range of legal `bigdatetime` values is from January 1, 0001 00:00:00.000000 to December 31, 9999 23:59:59.999999.
- `bigtime` indicates the number of microseconds that have passed since the beginning of the day. The range of legal `bigtime` values is from 00:00:00.000000 to 23:59:59.999999.
- `date` can hold dates from January 1, 0001 to December 31, 9999, exactly matching the allowable values in `java.sql.Date`. A direct mapping exists between `java.sql.Date` and the `date` datatype.
- `time` can hold time between 00:00:00.000 and 23:59:59.990. A direct mapping exists between `java.sql.Time` and the `time` datatype.

### Related Information

[Date, Time, Datetime, and Smalldatetime Datatypes \[page 78\]](#)

[Bigdatetime and Bigtime Datatypes \[page 79\]](#)

### 2.6.8.4.1 Date, Time, Datetime, and Smalldatetime Datatypes

SAP jConnect for JDBC supports `date`, `time`, `datetime`, and `smalldatetime`.

If you select from a table that contains a `date` or `time` column, and you have not enabled `date/time` support in SAP jConnect (by setting the version), the server tries to convert the `date` or `time` to a `datetime` value before returning it.

- This might cause problems if the date to be returned is earlier than 1/1/1753. In that case, a conversion error occurs, and the database informs you of the error.
- SAP SQL Anywhere supports a `date` and `time` datatype, but they are not yet directly compatible with those in SAP Adaptive Server Enterprise (SAP ASE) version 12.5.1 and later. Using SAP jConnect, continue to use the `datetime` and `smalldatetime` datatypes when communicating with SAP SQL Anywhere.

- The maximum value in a `datetime` column in SAP SQL Anywhere is 1-1-7911 00:00:00.
- Using SAP jConnect, you receive conversion errors if you attempt to insert dates earlier than 1/1/1753 into `datetime` columns or parameters.
- Refer to the SAP ASE manuals for more information on the `date` and `time` datatypes; of special note is the information about allowable implicit conversions.
- If you use `getObject` with an SAP ASE `date`, `time`, or `datetime` column, the value returned is, respectively, a `java.sql.Date`, `java.sql.Time`, or `java.sql.Timestamp` datatype.

## 2.6.8.4.2 Bigdatetime and Bigtime Datatypes

When connecting to SAP Adaptive Server Enterprise (SAP ASE) 15.5 and later, SAP jConnect transfers data using the `bigdatetime` and `bigtime` datatypes even if the receiving SAP ASE columns are defined as `datetime` and `time`.

- This means that SAP ASE may silently truncate the values from SAP jConnect to fit SAP ASE columns. For example, a `bigtime` value of 23:59:59.999999 is saved as 23:59:59.996 in an SAP ASE server column with datatype `time`.
- When connecting to SAP ASE 15.0.x and earlier, SAP jConnect for JDBC transfers data using the `datetime` and `time` datatypes.

## 2.6.8.5 Char, Varchar, Text, and GetByte Datatypes

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

## 2.6.8.6 Other Supported Datatypes

Review other SAP Adaptive Server Enterprise (SAP ASE) datatypes supported by SAP jConnect.

SAP jConnect supports these SAP ASE datatypes:

- `bigint` – an exact numeric datatype designed to be used when the range of the existing `int` types is insufficient.
- `unsigned int` – unsigned versions of the exact numeric integer datatypes: `unsignedsmallint`, `unsignedint`, and `unsignedbigint`.
- `unitext` – a variable-length datatype for Unicode characters.

## Related Information

[Bigint Datatype \[page 80\]](#)

[Unitext Datatype \[page 80\]](#)

[Unsigned Int Datatypes \[page 80\]](#)

## 2.6.8.6.1 Bigint Datatype

`bigint` which is a 64-bit integer datatype that is supported as a native SAP Adaptive Server Enterprise (SAP ASE) datatype.

`bigint` maps to the Java datatype `long`. To use this as a parameter, call `PreparedStatement.setLong(int index, long value)` and SAP jConnect sends the data as `bigint` to SAP ASE. When retrieving from a `bigint` column, use the `ResultSet.getLong(int index)` method.

## 2.6.8.6.2 Unitext Datatype

SAP jConnect internally stores and retrieves data from SAP Adaptive Server Enterprise (SAP ASE) when `unitext` columns are used.

## 2.6.8.6.3 Unsigned Int Datatypes

SAP Adaptive Server Enterprise (SAP ASE) supports `unsigned bigint`, `unsigned int`, and `unsigned smallint` as native SAP ASE datatypes.

Because, there are no corresponding unsigned datatypes in Java, you must `set` and `get` the next higher integer to process the data correctly. For example, if you are retrieving data from an `<unsigned int>`, using the Java datatype `int` is too small to contain positive large values, and as a result, `ResultSet.getInt(int index)` might return incorrect data or throw an exception. To process the data correctly, get the next higher integer value `ResultSet.getLong()`.

SAP ASE Datatype	Java Datatype
<code>unsigned smallint</code>	<code>setInt()</code> , <code>getInt()</code>
<code>unsigned int</code>	<code>setLong()</code> , <code>getLong()</code>
<code>unsigned bigint</code>	<code>setBigDecimal()</code> , <code>getBigDecimal()</code>

## 2.7 Variable-Length Rows in Data-Only-Locked Tables

Versions of SAP Adaptive Server Enterprise (SAP ASE) earlier than 15.7 configured for 16K logical page sizes could not create data-only-locked (DOL) tables with variable-length rows if a variable-length column began more than 8191 bytes after the start of the row.

This limitation has been removed starting in SAP ASE 15.7. See *Data Storage* in the *SAP ASE Enterprise Performance and Tuning Series: Physical Database Tuning*.

JDBC clients do not need special configuration to use this feature. When connected to SAP ASE version 15.7 that has been configured to receive wide DOL rows, these clients automatically insert records using the wide offset. An error message is received if a client attempts to send a wide DOL row to an earlier version of SAP ASE, or to SAP ASE 15.7 for which the wide DOL row option is disabled.

## 2.8 Large Object (LOB) Support

SAP jConnect supports using large object (LOB) datatypes—`text`, `untext`, and `image`

- LOB columns with in-row storage – In SAP Adaptive Server Enterprise (SAP ASE), LOB columns that are marked for in-row are stored in-row when there is adequate memory to hold the entire row. When the size of a row increases over its defined limit due to an update to any column in it, the LOB columns which are stored in-row are moved off-row to bring it within the limits. See *In-Row Off-Row LOB* in the *SAP ASE Enterprise Transact-SQL Users Guide*.

The bulk insert routines in SAP jConnect support the in-row and off-row storage of `text`, `image`, and `untext` LOB columns in SAP ASE. Bulk insert routines from earlier client versions always store LOB columns off-row.

- LOB objects as parameters of stored procedures – SAP jConnect supports using `text`, `untext`, and `image` as input parameters in stored procedures and as parameter marker datatypes.

## 2.9 Large Object Locator Support

SAP jConnect supports large object (LOB) locators. A LOB locator contains a logical pointer to LOB data rather than the data itself, reducing the amount of data that passes through the network between SAP Adaptive Server Enterprise (SAP ASE) and its clients.

Server support for LOB locators was introduced in SAP ASE 15.7.

SAP jConnect accesses LOB data using server-side locators when connected to an SAP ASE that supports LOB locators and `autocommit` is turned off. Otherwise, jConnect materializes LOB data at the client side. You can use the complete LOB API with client-side materialized LOB data, however, due to larger data, API performance may be different than when used with LOB locators.

### i Note

When you are using LOB locators, retrieving a large result set that includes LOB data on each row may impact your application's performance. SAP ASE returns a LOB locator as part of the result set and, to

obtain LOB data, SAP jConnect must cache the remaining result set. SAP recommends that you keep result sets small, or that you enable cursor support to limit the size of data to be cached.

To enable LOB locator support, establish a connection to SAP ASE with the `ENABLE_LOB_LOCATORS` connection property set to true. When enabled, client applications can access the locators using the `Blob`, `Clob`, and `NClob` classes from the `java.sql` package.

### **i Note**

When both LOB locators and `autocommit` are enabled, SAP jConnect automatically switches the LOB locators to client-side-materialized LOB even if the SAP ASE is capable of supporting LOB locators. This increases the memory used by the client and may degrade performance. Therefore, it is advisable to use LOB locators with `autocommit off`.

For information about the `Blob`, `Clob`, and `NClob` classes, see the Java documentation.

## **2.10 Advanced Features in SAP jConnect**

SAP jConnect provides advanced features such as event notification, error message handling, password encryption, dynamic class loading, and JDBC specification support.

Review the instructions to use the advanced features supported by SAP jConnect.

### **Related Information**

[BCP Insert \[page 83\]](#)  
[Supported SAP ASE Cluster Edition Features \[page 84\]](#)  
[Event Notification \[page 86\]](#)  
[Error Messages \[page 88\]](#)  
[Password Encryption \[page 92\]](#)  
[Store Java Objects as Column Data in Table \[page 95\]](#)  
[Dynamic Class Loading \[page 100\]](#)  
[JDBC 4.0 Specifications Support \[page 103\]](#)  
[JDBC 3.0 Specifications Support \[page 104\]](#)  
[Support for JDBC 2.0 Optional Package Extensions \[page 106\]](#)  
[On-demand Command Encryption for SAP ASE Clients \[page 115\]](#)  
[Multiversion Concurrency Control \[page 118\]](#)  
[BCP Insert \[page 83\]](#)  
[Supported SAP ASE Cluster Edition Features \[page 84\]](#)  
[Event Notification \[page 86\]](#)  
[Error Messages \[page 88\]](#)  
[Password Encryption \[page 92\]](#)

[JDBC 4.0 Specifications Support \[page 103\]](#)

[Store Java Objects as Column Data in Table \[page 95\]](#)

[Dynamic Class Loading \[page 100\]](#)

[JDBC 3.0 Specifications Support \[page 104\]](#)

[Support for JDBC 2.0 Optional Package Extensions \[page 106\]](#)

## 2.10.1 BCP Insert

SAP jConnect supports large insertions of rows to SAP Adaptive Server Enterprise (SAP ASE) version 12.5.2 and later using bulk-load inserts.

Although this feature does not require special server configuration, a larger page size, network packet size, and maximum memory size significantly improves performance.

Additionally, depending on the client memory, using of larger batches also improves performance.

To enable bulk-load inserts, set `ENABLE_BULK_LOAD` to one of these values:

- `ARRAYINSERT_WITH_MIXED_STATEMENTS` – enables bulk loading with row-level logging, and allows your application to execute other statements during the bulk-load operation.
- `ARRAYINSERT` – enables bulk loading with row-level logging, but your application cannot execute other statements during the bulk-load operation.
- `BCP` – enables bulk loading with page-level logging; your application cannot execute other statements during the bulk-load operation.
- `LOG_BCP` – enables the bulk loading with page-level logging using the SAP ASE fast-log BCP feature; your application cannot execute other statements during the bulk-load operation.

When you use prepared statements, and `ENABLE_BULK_LOAD` is set to a valid value, SAP jConnect uses the BULK routines to insert a batch of records into SAP databases.

## Limitations for Enabling Bulk Load

Some of the limitations for `ENABLE_BULK_LOAD`:

- Triggers are ignored on selected tables.
- Null and referential constraints are not validated.
- Computed and encrypted columns are not supported.
- Duplicate identity values might be created if you specify identity value ranges that overlap.
- Conflicting `IDENTITY` values might be inserted when concurrent data entry occurs.

### i Note

`ENABLE_BULK_LOAD` property is specifically added to enhance the performance of insert operation. When a connection property is set, data is sent in bulk; therefore, driver doesn't get the information on the update count for each row, instead it receives the cumulative count of total rows affected in a batch.

When the `ENABLE_BULK_LOAD` property is set, `Statement/PerparedStatement/CallableStatement.executeBatch()` returns the total affected rows instead of the status of each row.

## 2.10.2 Supported SAP ASE Cluster Edition Features

SAP jConnect supports the SAP Adaptive Server Enterprise (SAP ASE) Cluster Edition environment, where multiple SAP ASE servers connect to a shared set of disks and a high-speed private interconnection. This allows SAP ASE to scale using multiple physical and logical hosts.

For more information about Cluster Edition, see the *SAP ASE Enterprise Users Guide to Clusters*.

### Related Information

[Login Redirection \[page 84\]](#)

[Connection Migration \[page 84\]](#)

[Connection Failover \[page 85\]](#)

### 2.10.2.1 Login Redirection

When a client application attempts to connect to a busy server, login redirection helps balance the load of the servers by allowing the server to redirect the client connection to less busy servers within the cluster.

At any given time, some servers within a cluster environment are usually more loaded with work than others. The login redirection occurs during the login sequence and the client application does not receive notification that it was redirected. Login redirection is enabled automatically when a client application connects to a server that supports this feature.

#### i Note

When a client application connects to a server that is configured to redirect clients, the login time may increase because the login process is restarted whenever a client connection is redirected to another server.

### 2.10.2.2 Connection Migration

Connection migration allows a server in a cluster environment to dynamically distribute load, and seamlessly migrate an existing client connection and its context to another server within the cluster.

This feature enables a cluster environment to achieve optimal resource utilization and decrease computing time. Because migration between servers is seamless, connection migration also helps create a highly available, zero-downtime environment. Connection migration is enabled automatically when a client application connects to a server that supports this feature.

#### i Note

Command execution time may increase during server migration. SAP recommends that you increase the command timeouts accordingly.

## 2.10.2.3 Connection Failover

Connection failover allows a client application to switch to an alternate SAP Adaptive Server Enterprise (SAP ASE) if the primary server becomes unavailable due to an unplanned event, like power outage or a socket failure.

In a cluster environment, client applications can fail over numerous times to multiple servers using dynamic failover addresses.

With high availability enabled, the client application does not need to be configured to know the possible failover targets. SAP ASE keeps the client updated with the best failover list based on cluster membership, logical cluster usage, and load distribution. During failover, the client refers to the ordered failover list while attempting to reconnect. If the driver successfully connects to a server, the driver internally updates the list of host values based on the list returned. Otherwise, the driver throws a connection failure exception.

### i Note

The connection properties `DEFAULT_QUERY_TIMEOUT` and `INTERNAL_QUERY_TIMEOUT` or `DriverManager.setLoginTimeout(xx)` play vital role to switch over from failed node to highly available node after failover occurs.

## Related Information

[Enabling Connection Failover \[page 85\]](#)

### 2.10.2.3.1 Enabling Connection Failover

You can use the connection string to enable connection failover by setting `REQUEST_HA_SESSION` to true.

### ❖ Example

For example:

```
URL="jdbc:sybase:Tds:server1:port1,server2:port2,...,
serverN:portN/mydb?REQUEST_HA_SESSION=true"
```

where `server1:port1, server2:port2, ... , serverN:portN` is the ordered failover list.

SAP jConnect attempts to connect to the first host and port specified in the failover list. If unsuccessful, goes through the list until a connection is established or until the end of the list is reached.

### i Note

The list of alternate servers specified in the connection string is used only during initial connection. After the connection is established with any available instance, and if the client supports high availability, the client receives an updated list of the best possible failover targets from the server. This new list overrides the specified list.



## 2.10.3 Event Notification

You can use event notification to have your application notified when an SAP 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`, which 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:

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

where:

- `<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 SAP Open Server, SAP 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 SAP Open Server cannot add the client to the notification list, the call to `regWatch` throws a SQL exception.

To turn off event notification:

```
SybConnection.regNoWatch(<proc_name>)
```

### ⚠ Caution

When you use SAP jConnect event notification extensions, the application must 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 stop responding when it exits the application.

## Related Information

[Event Notification Example \[page 87\]](#)

## 2.10.3.1 Event Notification Example

Review instructions to implement an event handler and then register an event with an instance for your event handler, once a connection is established.

Event notification sample code:

```
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();
```

```
}
```

## 2.10.4 Error Messages

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

### Related Information

[Numeric Errors Returned as Warnings \[page 88\]](#)

[Retrieve SAP jConnect-Specific Error Information \[page 89\]](#)

[Customizing Error-Message Handling \[page 90\]](#)

[Installing an Error-Message Handler \[page 91\]](#)

[Error Message Handler Example \[page 91\]](#)

### 2.10.4.1 Numeric Errors Returned as Warnings

Numeric errors are handled by default as severity 10 in SAP Adaptive Server Enterprise (SAP ASE) 12.0 through 12.5.

A severity-level 10 message is classified as a status information message, not as an error, and its content is transferred in a `SQLWarning` object.

This code illustrates the process:

```
static void processWarnings(SQLWarning warning)
{
    if (warning != null)
    {
        System.out.println ("\n -- Warning received -- \n");
    } //end if
    while (warning != null)
    {
        System.out.println ("Message: " + warning.getMessage());
        System.out.println ("SQLState: " + warning.getSQLState());
        System.out.println ("ErrorCode: " + warning.getErrorCode());
        System.out.println ("-----");
        warning = warning.getNextWarning();
    } //end while
} //end processWarnings
```

When a numeric error occurs, the `ResultSet` object returned contains no result set data, and the relevant information concerning the error must be obtained from the `SQLWarning`. Therefore, in a JDBC application, the

code that checks for and processes a `SQLWarning` should not depend on a result set. For example, the following code checks for and processes `SQLWarning` data both inside and outside the result-set processing while loop:

```
while (rs.next())
{
    String value = rs.getString(1);
    System.out.println ("Fetched value: " + value);
    // Check for SQLWarning on the result set.
    processWarnings (rs.getWarnings());
} //end while
// Check for SQLWarning on the result set.
processWarnings (rs.getWarnings());
```

Here, the code checks for `SQLWarning` even if there is no result set data (`rs.next()` is false). The following example is output for a program properly written to detect and report numeric errors. The error is a division by zero:

```
-- Warning received --
Message: Divide by zero occurred.
SQLState: 01012
ErrorCode: 3607
```

## 2.10.4.2 Retrieve SAP jConnect-Specific Error Information

SAP jConnect provides an `EedInfo` interface that specifies methods for obtaining SAP jConnect-specific error information.

The `EedInfo` interface is implemented in `SybSQLException` and `SybSQLWarning`, which extend the `SQLException` and `SQLWarning` classes.

`SybSQLException` and `SybSQLWarning` contain these methods:

- `public ResultSet getEedParams`, which returns a one-row result set containing any parameter values that accompany the error message.
- `public int getStatus`, which returns a 1 if there are parameter values, and returns a 0 if there are no parameter values in the message.
- `public int getLineNumber`, which returns the line number of the stored procedure or query that caused the error message.
- `public String getProcedureName`, which returns the name of the procedure that caused the error message.
- `public String getServerName`, which returns the name of the server that generated the message.
- `public int getSeverity`, which returns the severity of the error message.
- `public int getState`, which returns information about the internal source of the error message in the server for use only by SAP Technical Support.
- `public int getTranState`, which 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 can be `SQLException` or `SQLWarning` messages without being `SySQLException` or `SybsQLWarning` messages. Your application should check the type of exception it is handling before it downcasts to `SySQLException` or `SybsQLWarning`.

## 2.10.4.3 Customizing Error-Message Handling

Use the `SybMessageHandler` interface to customize the way SAP jConnect handles error messages generated by the server.

### Context

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 severity rating of the server. 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. However, you may want to upgrade the severity in other circumstances.

#### i Note

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

#### ❖ Example

When SAP jConnect receives an error message, it checks to see if a `SybMessageHandler` class has been registered for handling the message. If so, SAP jConnect invokes the `messageHandler` method, which accepts a SQL exception as its argument. SAP jConnect then 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, SAP 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 once the control is returned to SAP jConnect.

## 2.10.4.4 Installing an Error-Message Handler

Install an error-message handler implementing `SybMessageHandler` by calling the `setMessageHandler` method from `SybDriver`, `SybConnection`, or `SybStatement`.

### Context

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`.

#### ❖ Example

This hierarchy only applies from the time the error-message handler object is installed. For example, if you create a `SybConnection` object called "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`.

## 2.10.4.5 Error Message Handler Example

Example for an error message handler in SAP jConnect.

```
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.jdbc4.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'");
        }
        catch (Exception e)
        {
            e.printStackTrace();
        }
    }
}
```

```

        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();
}

```

## 2.10.5 Password Encryption

By default, SAP jConnect for JDBC sends plain text passwords over the network to SAP Adaptive Server Enterprise (SAP ASE) for authentication.

However, SAP jConnect also supports symmetrical and asymmetrical password encryption and can encrypt passwords before they are sent over the network.

The symmetrical encryption mechanism uses the same key to encrypt and decrypt the password, whereas an asymmetrical encryption mechanism uses one key (the public key) to encrypt the password and another key (the private key) to decrypt the password. Because the private key is not shared across the network, asymmetrical encryption is considered more secure than symmetrical encryption. When password encryption is enabled, and the server supports asymmetric encryption, this format is used instead of symmetric encryption.

### **i Note**

To use the asymmetric password encryption feature, you must have a server that supports password encryption, such as SAP ASE 15.0.2 or later.

## **Related Information**

[Enabling Password Encryption \[page 93\]](#)

[Enabling Login Retry with Clear Text Password \[page 93\]](#)

[Setting Up the Java Cryptography Extension \(JCE\) Provider \[page 94\]](#)

## **2.10.5.1 Enabling Password Encryption**

The `ENCRYPT_PASSWORD` connection property specifies whether the password is transmitted in encrypted format.

### **Context**

This same property enables asymmetric key encryption. When password encryption is enabled and the server supports asymmetric key encryption, this format is used instead of symmetric key encryption.

Set the `ENCRYPT_PASSWORD` connection property to true to enable password encryption. The default value is false.

### **i Note**

If the server is configured to require clients to use an encrypted password, entering a plain text password causes user login to fail.

## **2.10.5.2 Enabling Login Retry with Clear Text Password**

Server login fails when the `ENCRYPT_PASSWORD` property is set to true, and the server does not support password encryption.

### **Context**

To use a clear text password for servers that do not support password encryption, set the `RETRY_WITH_NO_ENCRYPTION` connection property to true.



When both `ENCRYPT_PASSWORD` and `RETRY_WITH_NO_ENCRYPTION` properties are set to true, jConnect first logs in using the encrypted password. If login fails, jConnect logs in using the clear text password.

### 2.10.5.3 Setting Up the Java Cryptography Extension (JCE) Provider

Asymmetric password encryption mechanism uses RSA encryption algorithms to encrypt the password being transmitted.

#### Context

To perform this RSA encryption, configure your JRE with a suitable Java Cryptography Extension (JCE) provider. The configured JCE provider should be capable of supporting the "RSA/ECB/OAEPWithSHA1AndMGF1Padding" or "RSA/NONE/OAEPWithSHA1AndMGF1Padding" transformation. SAP jConnect ships a FIPS 140-2 certified JCE Provider with the product and it is used by default.

#### ❖ Example

You can use the `JCE_PROVIDER_CLASS` connection property to specify an alternate JCE provider. There are a number of commercial and open source JCE providers that you can choose from. For example, the "Bouncy Castle Crypto APIs for Java" is a popular open source Java JCE provider. If you choose not to specify the `JCE_PROVIDER_CLASS` property, SAP jConnect uses the default FIPS 140-2 certified bundled provider.

#### Related Information

[Using Default JCE Provider to Perform Password Encryption \[page 95\]](#)

[Specifying Custom JCE Provider \[page 95\]](#)

### 2.10.5.3.1 Using Default JCE Provider to Perform Password Encryption

SAP jConnect uses default bundled FIPS 140-2 certified JCE Provider to perform RSA password encryption, when `JCE_PROVIDER_CLASS` is not specified.

### 2.10.5.3.2 Specifying Custom JCE Provider

Specify a custom JCE provider in SAP jConnect.

#### Procedure

1. Set the `JCE_PROVIDER_CLASS` property to the fully qualified class name of the provider you want to use.

For example, to use the Bouncy Castle JCE:

```
String url = "jdbc:sybase:Tds:myserver:3697";
Properties props = new Properties();
props.put("ENCRYPT_PASSWORD ", "true");
props.put("JCE_PROVIDER_CLASS",
"org.bouncycastle.jce.provider.BouncyCastleProvider");
/* Set up additional connection properties as needed */
props.put("user", "xyz");
props.put("password", "123");
/* get the connection */
Connection con = DriverManager.getConnection(url, props);
```

2. Configure the JCE provider.

Either:

- Copy the JCE provider `jar` file into the JRE standard extension directory:
  - For UNIX platforms: `${JAVA_HOME}/jre/lib/ext`
  - For Windows: `%JAVA_HOME%\jre\lib\ext`
- Or, if you cannot copy the JCE `jar` file to the appropriate directory, see *JCE Reference Guide* for instructions on setting up an external JCE provider.

If no other JCE providers are configured, or if configured providers do not support the required transformation and password encryption is enabled, the connection fails.

## 2.10.6 Store Java Objects as Column Data in Table

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.

SAP 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 SAP jConnect with an application that uses native JDBC classes and methods to directly store and retrieve Java objects as column data.

### **i Note**

To use `getObject` and `setObject`, set the SAP jConnect version to `com.sybase.jdbcx.SybDriver.VERSION_4` or later.

SAP Adaptive Server Enterprise (SAP ASE) version 12.0 and later, and SAP SQL Anywhere version 6.0.x and later can store Java objects in a table, with some limitations. See the *SAP jConnect for JDBC Release Bulletin*.

## **Related Information**

[Prerequisites for Storing Java Objects as Column Data \[page 96\]](#)

[Sending Java Objects to Database \[page 97\]](#)

[Receive Java Objects from Database \[page 98\]](#)

[Prerequisites for Storing Java Objects as Column Data \[page 96\]](#)

[Receive Java Objects from Database \[page 98\]](#)

[JCONNECT\\_VERSION Connection Property \[page 9\]](#)

[Sending Java Objects to Database \[page 97\]](#)

[BCP Insert \[page 83\]](#)

[Supported SAP ASE Cluster Edition Features \[page 84\]](#)

[Event Notification \[page 86\]](#)

[Error Messages \[page 88\]](#)

[Password Encryption \[page 92\]](#)

[JDBC 4.0 Specifications Support \[page 103\]](#)

## **2.10.6.1 Prerequisites for Storing Java Objects as Column Data**

Store Java objects belonging to a user-defined Java class in a column.

- The class must implement the `java.io.Serializable` interface. This is because SAP 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 SAP SQL Anywhere or an SAP Adaptive Server Enterprise (SAP ASE) server and use it as if it were present in the local CLASSPATH.
- The client system must have the class definition in a `.class` file that is accessible through the local CLASSPATH environment variable.

## Related Information

[Dynamic Class Loading \[page 100\]](#)

### 2.10.6.2 Sending Java Objects to Database

Use the `setObject` methods to send Java objects to a database.

#### Context

To send an instance of a user-defined class as column data, use one of the `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`, to send a Java object, you can use the `java.sql.Types.JAVA_OBJECT` target `sql.Type`, or you can use `java.sql.Types.OTHER`.

#### ❖ Example

This 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();

```

### 2.10.6.3 Receive Java Objects from Database

Client JDBC applications 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 `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 this `getObject` method in the `CallableStatement` interface to retrieve the object:

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

This 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 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 SAP jConnect installation directory.

## 2.10.7 Dynamic Class Loading

SAP SQL Anywhere and SAP Adaptive Server Enterprise (SAP ASE) allow you to specify Java classes.

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

SAP jConnect version 6.05 and later implements `DynamicClassLoader` (DCL) to load a class directly from an SAP SQL Anywhere or SAP ASE server and use it as if it were present in the local CLASSPATH.

In SAP jConnect 6.0 and earlier versions, only classes that appeared in the SAP jConnect CLASSPATH were accessible, that is, any attempt of a SAP jConnect application to access an instance of a class that was not in the local CLASSPATH, resulted in a `java.lang.ClassNotFoundException`.

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

See *Java in Adaptive Server* for more detailed information about using Java and SAP ASE.

### Related Information

[Using DynamicClassLoader \[page 100\]](#)

[Deserialization \[page 102\]](#)

[Preloading .jar Files \[page 102\]](#)

[Additional Dynamic Class Loading Features \[page 103\]](#)

[BCP Insert \[page 83\]](#)

[Supported SAP ASE Cluster Edition Features \[page 84\]](#)

[Event Notification \[page 86\]](#)

[Error Messages \[page 88\]](#)

[Password Encryption \[page 92\]](#)

[JDBC 4.0 Specifications Support \[page 103\]](#)

### 2.10.7.1 Using DynamicClassLoader

Use the `CLASS_LOADER` connection property to provide a convenient mechanism for sharing one class loader among several connections.

#### Procedure

1. Create and configure a class loader.

The code for your SAP jConnect application 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 (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);
```

Assume the Java class definition is:

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

Assume the SQL table definition:

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

3. Use the client-side code in the absence of an `Addr` class in the client application 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.
    ...
}
```

## Results

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 might 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. Consequently, the second class is never loaded, and there is no direct way for SAP jConnect 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 or server.

## 2.10.7.2 Deserialization

The serialized object is an instance of a class that resides on a server and does not exist in the CLASSPATH.

This example illustrates how to deserialize an object from a local file:

`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();
```

## 2.10.7.3 Preloading .jar Files

SAP jConnect version 6.05 or later has 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.

### Context

In this context, "JAR" refers to the "retained JARname" used by the server. This is the `.jar` file 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` file data is retrieved from the server unnecessarily.

#### **i Note**

SAP SQL Anywhere cannot return a `.jar` file as one entity, so SAP jConnect iteratively retrieves each class in turn. However, SAP ASE retrieves the entire `.jar` file and loads each class that it contains.

## **2.10.7.4 Additional Dynamic Class Loading Features**

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

You can use public methods inherited from `java.lang.ClassLoader`. Methods in `java.lang.Class` that deal with loading classes are also available; however, use these methods with caution because some of them make assumptions about which class loader gets used. In particular, you should use the three-argument version of `Class.forName`, otherwise the system (boot) class loader is used.

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

## **Related Information**

[Error Messages \[page 88\]](#)

## **2.10.8 JDBC 4.0 Specifications Support**

Some JDBC 4.0 specifications that are supported by SAP jConnect.

- Connection management
- Automatic SQL driver loading
- Database metadata
- National character set conversion
- Wrapper pattern
- Scalar functions `CHAR_LENGTH`, `CHARACTER_LENGTH`, `CURRENT_DATE`, `CURRENT_TIME`, `CURRENT_TIMESTAMP`, `EXTRACT`, and `OCTET_LENGTH`, `POSITION`

See [Oracle Technology Network for Java](#) 🐘 for information about the JDBC 4.0 specifications.

## 2.10.9 JDBC 3.0 Specifications Support

JDBC 3.0 features that are supported in SAP jConnect 16.0.

### Related Information

[Savepoint Support \[page 104\]](#)  
[Retrieval of Parameter Metadata \[page 105\]](#)  
[Retrieval of Autogenerated Keys \[page 105\]](#)  
[Multiple Open ResultSet Objects \[page 105\]](#)  
[Pass Parameters to CallableStatement Objects by Name \[page 106\]](#)  
[Holdable Cursor Support \[page 106\]](#)  
[BCP Insert \[page 83\]](#)  
[Supported SAP ASE Cluster Edition Features \[page 84\]](#)  
[Event Notification \[page 86\]](#)  
[Error Messages \[page 88\]](#)  
[Password Encryption \[page 92\]](#)  
[JDBC 4.0 Specifications Support \[page 103\]](#)

### 2.10.9.1 Savepoint Support

You can use the `Savepoint` interface, which contains methods to set, release, or roll back a transaction to designated savepoints.

<b>Using Savepoints in your transactions</b>	JDBC 3.0, the <code>Savepoint</code> interface allows you to partition a transaction into logical breakpoints, providing control over how much of the transaction gets rolled back.
<b>Setting and rolling back to a Savepoint</b>	JDBC 3.0 API includes the method <code>Connection.setSavepoint</code> , which sets a savepoint within the current transaction and returns a <code>Savepoint</code> object. The <code>Connection.rollback</code> method is overloaded to take a <code>Savepoint</code> object argument.
<b>Releasing a Savepoint</b>	<p>The <code>Connection.releaseSavepoint</code> method takes a <code>Savepoint</code> object as a parameter and removes it from the current transaction.</p> <p>After a <code>Savepoint</code> has been released, if you try to reference it in a rollback operation, a <code>SQLException</code> occurs. Any savepoints you create in a transaction are automatically released and become invalid when the transaction is committed or when the entire transaction is rolled back. If you roll a transaction back to a savepoint, it automatically releases and invalidates any other savepoints that were created after the savepoint in question.</p>

### i Note

You can use the `DatabaseMetaData.supportsSavepoints` method to determine whether a JDBC API implementation supports savepoints.

## 2.10.9.2 Retrieval of Parameter Metadata

The JDBC 3.0 `ParameterMetaData` interface describes the number, type, and properties of parameters to prepared statements, and supports the most current `DatabaseMetaData` methods.

## 2.10.9.3 Retrieval of Autogenerated Keys

JDBC 3.0 addresses the common need to obtain from columns the value of an autogenerated or autoincremented key.

To retrieve autogenerated keys, pass the constant `Statement.RETURN_GENERATED_KEYS` as the second parameter of the `Statement.execute()` method.

After you have executed the statement, call `Statement.getGeneratedKeys()` to retrieve the generated keys. The result set contains a row for each retrieved key.

### i Note

SAP Adaptive Server Enterprise (SAP ASE) cannot return a result set of generated keys. If you execute a batch of `insert` commands, invoking `Statement.getGeneratedKeys()` returns only the value of the last generated key.

For more information about retrieving autogenerated keys, including a sample code, search for “retrieving automatically generated keys” on the Oracle Java Web site.

## 2.10.9.4 Multiple Open ResultSet Objects

JDBC 3.0 includes `getMoreResults(int)`, which takes an argument that specifies whether `ResultSet` objects returned by a `Statement` object should be closed before returning any subsequent `ResultSet` objects.

The JDBC 3.0 specification allows the `Statement` interface to support multiple open `ResultSets`, which removes the limitation of the JDBC 2.0 specification that statements returning multiple results must have only one `ResultSet` open at any given time. To support multiple open results, the `Statement` interface adds an overloaded version of the method `getMoreResults()`. The `getMoreResults(int)` method takes an integer flag that specifies the behavior of previously opened `ResultSets` when the `getResultSet()` method is called. The interface defines the flags as follows:

- `CLOSE_ALL_RESULTS` – all previously opened `ResultSet` objects are closed when calling `getMoreResults()`.
- `CLOSE_CURRENT_RESULT` – the current `ResultSet` object are closed when calling `getMoreResults()`.
- `KEEP_CURRENT_RESULT` – the current `ResultSet` object is not closed when calling `getMoreResults()`.

### 2.10.9.5 Pass Parameters to CallableStatement Objects by Name

Allows a string to identify the parameter to be set for a `CallableStatement` object.

You can use the `CallableStatement` interface to specify parameters by their names, rather than by parameter's index. This is useful when a procedure has many parameters with default values. You can use named parameters to specify only the values that have no default value.

### 2.10.9.6 Holdable Cursor Support

A holdable cursor, or result does not automatically close when the transaction that contains the cursor is committed. You must specify the holdability of a `ResultSet` object.

JDBC 3.0 supports specifying cursor holdability. You must specify the holdability of your `ResultSet` when you prepare a statement using the `createStatement()`, `prepareStatement()`, or `prepareCall()` methods. The holdability may be one of the following constants:

- `HOLD_CURSORS_OVER_COMMIT` – `ResultSet` objects (cursors) are not closed; they are held open when a `commit` operation is implicitly or explicitly performed.
- `CLOSE_CURSORS_AT_COMMIT` – `ResultSet` objects (cursors) are closed when a `commit` operation is implicitly or explicitly performed.

Closing a cursor when a transaction is committed usually results in better performance. Unless you require the cursor after the transaction, SAP recommends that you close the cursor when the `commit` operation is carried out. Because the specification does not define the default holdability of a `ResultSet`, its behavior depends on the implementation.

### 2.10.10 Support for JDBC 2.0 Optional Package Extensions

The JDBC 2.0 Optional Package (formerly the JDBC 2.0 Standard Extension API) defined several features that JDBC 2.0 drivers could implement.

SAP jConnect version 6.05 and later have implemented several of these optional package extension features:

- JNDI for naming conventions – works with any SAP DBMS supported by SAP jConnect
- Connection pooling – works with any SAP DBMS supported by SAP jConnect
- Distributed transaction management support – works only with SAP ASE

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

## Related Information

[JNDI for Naming Databases \[page 107\]](#)

[Connection Pooling \[page 111\]](#)

[Distributed Transaction Management Support \[page 113\]](#)

[BCP Insert \[page 83\]](#)

[Supported SAP ASE Cluster Edition Features \[page 84\]](#)

[Event Notification \[page 86\]](#)

[Error Messages \[page 88\]](#)

[Password Encryption \[page 92\]](#)

[JDBC 4.0 Specifications Support \[page 103\]](#)

## 2.10.10.1 JNDI for Naming Databases

Review the information for JNDI for naming databases.

## Related Information

[Reference \[page 107\]](#)

[Related Interfaces \[page 107\]](#)

[Usage \[page 108\]](#)

### 2.10.10.1.1 Reference

The JDBC 2.0 Optional Package (formerly the JDBC 2.0 Standard Extension API).

### 2.10.10.1.2 Related Interfaces

Related interfaces provide JDBC clients with an alternative to the standard approach for obtaining database connections.

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

- `javax.naming.spi.ObjectFactory`

Instead of invoking `Class.forName("com.sybase.jdbc4.jdbc.SybDriver")`, then passing a JDBC URL to the `DriverManager.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 SAP implementation of the `DataSource` object is `com.sybase.jdbcx.SybDataSource` (see the Javadoc for details). This implementation supports the standard properties using the design pattern for JavaBean components:

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

#### i Note

`roleName` is not supported.

SAP 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.jdbc4.SybObjectFactory`.

## 2.10.10.1.3 Usage

`DataSource` is used in different ways, in different applications.

All options are presented with some code examples. For more information, see the *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*), and the JNDI documentation on the Oracle Java Web site.

## Related Information

[Configuration by Administrator: LDAP \[page 109\]](#)

[Access by Client \[page 109\]](#)

[Programmatic Configuration \[page 110\]](#)

[Retrieve Datasource Object by Client \[page 110\]](#)

### 2.10.10.1.3.1 Configuration by Administrator: LDAP

SAP 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
```

### 2.10.10.1.3.2 Access by Client

A JDBC client application allows you to access the server name to obtain a reference to a `DataSource` object, instead of accessing the `DriverManager` and providing a JDBC URL.

This is a typical JDBC client application. Once you obtain the connection, the client code is identical to any other JDBC client code. The code is generic and references SAP jConnect only when setting the object factory property, which you can do as part of the environment setup.

The SAP 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.jdbc4.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
...
```

You need not explicitly pass the `Properties` to the `InitialContext` constructor 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.jdbc4.jdbc.SybObjectFactory
```



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

### 2.10.10.1.3.3 Programmatic Configuration

The purpose of programmatic configuration 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), the administrator runs this configuration utility (outlined as follows) and reassigns the logical name to the new data source configuration. This phase is typically done by the person who performs database system administration or application integration for their company. 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.jdbc4.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("java:comp/env/jdbc/myASE", ds);
```

Once you set up your `DataSource`, 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.

### 2.10.10.1.3.4 Retrieve Datasource Object 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("java:comp/env/jdbc/myASE");  
Connection conn = ds.getConnection();
```

## 2.10.10.2 Connection Pooling

Review the connection pooling instructions in SAP jConnect.

### Related Information

[Reference \[page 111\]](#)

[Related Interfaces \[page 111\]](#)

[Overview \[page 111\]](#)

### 2.10.10.2.1 Reference

Review the JDBC 2.0 Optional Package (formerly the JDBC 2.0 Standard Extension API).

### 2.10.10.2.2 Related Interfaces

Review the related interfaces in JDBC.

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

### 2.10.10.2.3 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 when using the application.

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.

The `com.sybase.jdbc4.jdbc.ConnectionPoolDataSource` class is provided to interact with connection pooling implementations. When you use `ConnectionPoolDataSource`, pool implementations listen to the

`PooledConnection`. The implementation is notified when you close the connection, or if you have 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. Transaction determines whether 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 the way 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 in the *JDBC 2.0 Optional Package* (formerly the *JDBC 2.0 Standard Extension API*).

## Related Information

[Configuration by Administrator: LDAP \[page 113\]](#)

[Access by Middle-Tier Clients \[page 113\]](#)

## 2.10.10.2.3.1 Configuration by Administrator: LDAP

Configure the LDAP by entering an additional line to your LDIF entry.

In this 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
```

### Related Information

[JNDI for Naming Databases \[page 107\]](#)

[Configuration by Administrator: LDAP \[page 109\]](#)

## 2.10.10.2.3.2 Access by Middle-Tier Clients

Initializes three properties `INITIAL_CONTEXT_FACTORY`, `PROVIDER_URL`, and `OBJECT_FACTORIES` and retrieves a `ConnectionPoolDataSource` object.

For a more complete code example, see `sample2/SimpleConnectionPool.java`. The fundamental difference between access by client and middle-tier client is:

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

## 2.10.10.3 Distributed Transaction Management Support

Provides a standard Java API for performing distributed transactions with SAP Adaptive Server Enterprise (SAP ASE). This feature is designed for use in a large multitier environment.

### Related Information

[Reference \[page 114\]](#)

[Related Interfaces \[page 114\]](#)

[Background and System Requirements \[page 114\]](#)

## 2.10.10.3.1 Reference

The JDBC 2.0 Optional Package (formerly the JDBC 2.0 Standard Extension API).

## 2.10.10.3.2 Related Interfaces

Review the related interfaces in JDBC.


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

## 2.10.10.3.3 Background and System Requirements

Use the `dtm_tm_role` to enable the Distribute Transaction Management support.

- SAP jConnect must be communicating directly with the resource manager within SAP Adaptive Server Enterprise (SAP ASE) version 12.0 and later, and the installation must have Distributed Transaction Management support.
- Any user who wants to participate in a distributed transaction must be granted `dtm_tm_role`, or the transactions fail.
- To use distributed transactions, you must install the stored procedures in the `/sp` directory. Refer to *Installing Stored Procedures* in the *SAP jConnect for JDBC Installation Guide*.

Distributed Transaction Management Support with Version 12.x



## Related Information

[Configuration by Administrator LDAP \[page 115\]](#)

[Access by Middle-Tier Clients \[page 115\]](#)

### 2.10.10.3.3.1 Configuration by Administrator LDAP

Configure the LDAP by entering an additional line to your LDIF entry.

In this 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:XADDataSource
```

## Related Information

[JNDI for Naming Databases \[page 107\]](#)

[Configuration by Administrator: LDAP \[page 109\]](#)

### 2.10.10.3.3.2 Access by Middle-Tier Clients

Initializes three properties `INITIAL_CONTEXT_FACTORY`, `PROVIDER_URL`, and `OBJECT_FACTORIES`, and retrieves a `XADDataSource` object.

For example:

```
...
XADDataSource xads = (XADDataSource) ctx.lookup ("server=myASE");
XAConnection xaconn = xads.getXAConnection ();
```

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

```
...
XADDataSource xads = (XADDataSource) ctx.lookup("servername=myASE");
XAConnection xaconn = xads.getXAConnection("my_username", "my_password");
```

## 2.10.11 On-demand Command Encryption for SAP ASE Clients

On-demand encryption allows you to encrypt commands and system procedures that contain sensitive data. On-demand encryption is supported for Open Client and other SAP Adaptive Server Enterprise (SAP ASE) SDK drivers, such as jConnect, ADO.NET, and ODBC. The client applications use Common Crypto Library for cryptographic services, such as encryption and decryption.

The client applications can encrypt the command before sending it to the server. The server decrypts the command before processing it. The system returns an error if you run on-demand encryption on a client and server that is not running SAP ASE version 16.0 SP03 or later. On-demand encryption is silently ignored when

the connection is already secured with SSL or Kerberos (with message confidentiality mode ON). You can set encryption either at a command or connection-level. Configuring encryption on a connection lets you encrypt all client commands on that connection. To enable on-demand encryption, client-specific properties (at connection-level and command-level) are added.

For more information, see *SAP Open Server* and *SDK for SAP Adaptive Server Enterprise* documentation.

## Related Information

[jConnect Client and On-demand Command Encryption \[page 116\]](#)

[Enabling On-demand Command Encryption for jConnect Client \[page 116\]](#)

[Disabling On-demand Command Encryption for jConnect Client \[page 117\]](#)

[Using Unlimited Strength Java Cryptography Extension Policy Files \[page 118\]](#)

### 2.10.11.1 jConnect Client and On-demand Command Encryption

Client applications using jConnect JDBC driver can use the on-demand encryption feature to send sensitive SQL commands in encrypted form to the server that supports this feature.

### 2.10.11.2 Enabling On-demand Command Encryption for jConnect Client

You can enable on-demand encryption for jConnect JDBC client commands.

Use these options:

- Enable on-demand encryption at the command level using the new `ENCRYPT_COMMAND` attribute. For example:

#### Code Syntax

```
props.put("ENCRYPT_COMMAND", "false"); // Default is FALSE
connection = DriverManager.getConnection(url, props);
// Once connection is established, turn ON encryption
connection.setClientInfo("ENCRYPT_COMMAND", "true");
connection.createStatement(query).execute();
// Once command execution is done, turn OFF encryption
connection.setClientInfo("ENCRYPT_COMMAND", "false");
```

#### Note

If a server does not support command encryption and the client attempts to dynamically set the `ENCRYPT_COMMAND` attribute, the `setClientInfo()` API throws an exception (`SQLClientInfoException`).

- Enable on-demand encryption at the connection level by using the connection property, `ENCRYPT_COMMAND`. The `ENCRYPT_COMMAND` property lets you encrypt all client commands on the connection. The default value of the `ENCRYPT_COMMAND` property is `FALSE`.

#### Code Syntax

```
props.put("ENCRYPT_COMMAND", "true");
connection = DriverManager.getConnection(url, props);
```

You can also set or clear the connection property (`ENCRYPT_COMMAND`) dynamically on a specific connection using the `Connection.setClientInfo(String prop, String value)` API. The connection property, when set to `TRUE`, allows you to send all commands in an encrypted form using a symmetric session key. Whereas setting its value to `FALSE` ensures that the commands are not encrypted.

#### Note

If a client has already established SSL connection, any attempt to encrypt (that is, using the `ENCRYPT_COMMAND` attribute in command) is ignored. If a client application, with the connection property set as `ENCRYPT_COMMAND = true`, attempts to log on to a server that does not support on-demand encryption, the client access request is ignored by the server, and a logon failure is reported.

## 2.10.11.3 Disabling On-demand Command Encryption for jConnect Client

You can disable on-demand encryption for jConnect JDBC client commands.

Use these options:

- Disable on-demand encryption at the command level using `ENCRYPT_COMMAND`. For example:

#### Code Syntax

```
connection = DriverManager.getConnection(url, props);
connection.setClientInfo("ENCRYPT_COMMAND", "false");
```

#### Note

If a server does not support command encryption and the client attempts to dynamically set the `ENCRYPT_COMMAND` attribute, the `setClientInfo()` API throws an exception (`SQLClientInfoException`).

- Disable on-demand encryption at the connection level using `ENCRYPT_COMMAND`. The default value of the `ENCRYPT_COMMAND` property is `FALSE`.

#### Code Syntax

```
props.put("ENCRYPT_COMMAND", "false");
connection = DriverManager.getConnection(url, props);
```



## 2.10.11.4 Using Unlimited Strength Java Cryptography Extension Policy Files

The version of Java Cryptography Extension (JCE) policy files bundled in the JDK 6 environment allows you to use strong but limited cryptography. To use on-demand encryption for JDBC jConnect clients, download the Unlimited Strength Java Cryptography Extension Policy Files from the following location:

<http://www.oracle.com/technetwork/java/javase/downloads/jce-6-download-429243.html> 

The `jce_policy-6.zip` contains the following two files:

- `local_policy.jar`
- `US_export_policy.jar`

Add the files to `$JRE_HOME/lib/security` directory.

## 2.10.12 Multiversion Concurrency Control

Multiversion concurrency control (MVCC) enables concurrent read-write operations, by allowing the server to lock rows for writing in one session while granting access to unaltered versions of these rows in another session, so that reads never wait for writes, and writes never wait for reads.

SAP ASE offers two types of MVCC:

- Regular MVCC – uses in-memory row versioning, where all inserts, updates, and deletes are performed in the in-memory row storage.
- On-disk multiversion concurrency control (on-disk MVCC) – stores row versions in a temporary database specified when the database is created or altered.

The two types use slightly different parameters to create or alter the database. In addition to the standard isolation levels offered by SAP ASE, SAP ASE uses three isolation levels specifically to enable multiversion concurrency control, which are not part of the standard SQL:

- Statement snapshot
- Transaction snapshot
- Readonly statement snapshot

### Related Information

[Connection Support in Multiversion Concurrency Control \[page 119\]](#)

[JDBC Support \[page 120\]](#)

## 2.10.12.1 Connection Support in Multiversion Concurrency Control

`ISOLATION_MODE` is a connection-level property that is used to enable the MVCC feature by converting ANSI isolation levels to snapshot isolation levels for MVCC- and on-disk MVCC-enabled tables.

### Note

SAP ASE provides the following techniques to migrate applications to use snapshot isolation seamlessly, and with minimal application changes. Perform these changes on the application in conjunction with any necessary server-side changes. If you do not set the `ISOLATION_MODE` property, existing applications default to run at ANSI isolation levels.

You can set the `ISOLATION_MODE` connection property for the client and map ANSI isolation levels to snapshot isolation levels. When `ISOLATION_MODE` is set to `snapshot`, the isolation level maps as:

- Isolation level 0 and 1 – mapping to the `statement snapshot` isolation level
- Isolation level 2 and 3 – mapping to the `transaction snapshot` isolation level

Once the `ISOLATION_MODE` connection property is set, the server overrides the `set transaction isolation` command from the connection or system procedure, which avoids having to make application changes to set isolation levels. For example, if a stored procedure sets the isolation level to ANSI isolation levels 0, 1, 2, or 3, but when using MVCC-enabled tables and the `ISOLATION_MODE` property is set, the ANSI isolation level for the stored procedure is converted to the `statement snapshot` or `transaction snapshot` isolation levels described above.

`ISOLATION_MODE` can have these values:

- `default` – the applications use the isolation levels specified by the application.
- `snapshot` – the application uses the snapshot isolation levels listed in the following table when it specifies traditional ANSI isolation levels for MVCC-enabled tables.
- `readonly_snapshot` – the application uses the snapshot isolation levels listed in the following table when it specifies traditional ANSI isolation levels for MVCC-enabled tables.

Session-level Isolation Level	Connection Isolation Mode set to DEFAULT	Connection Isolation Mode set to SNAPSHOT	Connection Isolation Mode set to READONLY_SNAPSHOT
<code>read uncommitted</code>	<code>read uncommitted</code>	<code>statement snapshot</code>	<code>readonly statement snapshot</code>
<code>read committed</code>	<code>read committed</code>	<code>statement snapshot</code>	<code>readonly statement snapshot</code>
<code>repeatable read</code>	<code>repeatable read</code>	<code>transaction snapshot</code>	<code>transaction snapshot</code>
<code>serializable read</code>	<code>serializable read</code>	<code>transaction snapshot</code>	<code>transaction snapshot</code>

Session-level Isolation Level	Connection Isolation Mode set to DEFAULT	Connection Isolation Mode set to SNAPSHOT	Connection Isolation Mode set to READONLY_SNAPSHOT
statement snapshot	statement snapshot	statement snapshot	statement snapshot
transaction snapshot	transaction snapshot	transaction snapshot	transaction snapshot
readonly statement snapshot	readonly statement snapshot	readonly statement snapshot	readonly statement snapshot

## 2.10.12.2 JDBC Support

This example describes the JDBC support for multiversion concurrency control using `ISOLATION_MODE`:

- Property Name: `ISOLATION_MODE`
- Values: 0 / 1 / 2 (as described in the table above)
- 0: default
- 1: snapshot
- 2: readonly\_snapshot

At runtime, the client application can set the `ISOLATION_MODE` property using this API:

```
connection.setClientInfo("ISOLATION_MODE", "1");
connection.setClientInfo("ISOLATION_MODE", "2");
```

The code is:

```
public static void main(String args[])
{
    Properties props = new Properties();
    props.put("user", "USER");
    props.put("password", "PASSWORD");
    // Client application can used MVCC feature
    props.put("ISOLATION_MODE", "1");
    // 1 = SNAPSHOT
    // 2 = READONLY SNAPSHOT
    String url = "jdbc:sybase:Tds:HOST:PORT";
    Connection conn = null; conn = DriverManager.getConnection (url, props);
    // Dynamically changing isolation mode
    conn.setClientInfo("ISOLATION_MODE", "2");
    ...
    conn.setTransactionIsolation(SybConnection.TRANSACTION_STATEMENT_SNAPSHOT);
    ...
}
```

## 2.11 Restrictions and Interpretations of JDBC Standards

The SAP jConnect implementation of JDBC deviates from the JDBC standards.

### Related Information

[Unsupported JDBC 4.0 Specification Requirements \[page 121\]](#)

[Use Connection.isClosed and IS\\_CLOSED\\_TEST \[page 121\]](#)

[Statement.close with Unprocessed Results \[page 122\]](#)

[Adjustments for Multithreading \[page 123\]](#)

[ResultSet.getCursorName \[page 123\]](#)

[Execute Stored Procedures \[page 123\]](#)

[Adjustments for Multithreading \[page 123\]](#)

[Unsupported JDBC 4.0 Specification Requirements \[page 121\]](#)

[Use Connection.isClosed and IS\\_CLOSED\\_TEST \[page 121\]](#)

[Statement.close with Unprocessed Results \[page 122\]](#)

[ResultSet.getCursorName \[page 123\]](#)

[Execute Stored Procedures \[page 123\]](#)

### 2.11.1 Unsupported JDBC 4.0 Specification Requirements

Review the JDBC 4.0 statements that are not supported in this release.

- `java.sql.RowID`
- XML APIs introduced in JDBC 4.0

### 2.11.2 Use Connection.isClosed and IS\_CLOSED\_TEST

SAP jConnect offers a default interpretation of the `isClosed` method that differs from the behavior defined in the JDBC 4.0 specification.

When you call `Connection.isClosed`, SAP jConnect verifies that `Connection.close` has been called on this connection. If `close` has been called, SAP jConnect returns `true` for `isClosed`. However, if `Connection.close` has not been called, SAP jConnect tries to execute the `sp_mda` on the database. `sp_mda` is part of the standard metadata that SAP jConnect users must install when they use SAP jConnect with a database.

According to section 11.1 of the JDBC 4.0 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 whether a database connection is

valid. A typical client can determine that a connection is invalid by catching the exception that is thrown when an operation is attempted.

The purpose of calling `sp_mda` is so that SAP 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, SAP jConnect returns false for `isClosed` because it has verified that the database connection is valid and working. However, if the call to `sp_mda` results in a `SQLException` being thrown, SAP jConnect catches the exception and returns true for `isClosed` because it appears that there is something wrong with the connection.

To force SAP jConnect to more closely follow the standard JDBC behavior for `isClosed()`, set the `IS_CLOSED_TEST` connection property to the special value "INTERNAL." The INTERNAL setting means that SAP jConnect returns true for `isClosed` only when `Connection.close` has been called, or when SAP 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 intend for SAP jConnect to attempt a `select 1` when `isClosed` is called, set the `IS_CLOSED_TEST` connection property to `select 1`.

## 2.11.3 Statement.close with Unprocessed Results

The JDBC specification does not clearly address 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 `ResultSets`) returned by the `Statement`.

For example, assume that there is a stored procedure on the database that performs seven row inserts. An application then executes that stored procedure using a `Statement.execute`. In this case, a SAP database returns 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 SE documentation](#) in the Javadoc for the `java.sql.*` package.

However, an application programmer might incorrectly call `Statement.close` before reading through all of the returned update counts. In this case, SAP jConnect sends a `cancel` to the database, which might have unexpected and unwanted side effects.

In this particular example, if the application calls `Statement.close` before the database has completed the inserts, the database might not execute all of the inserts. It might stop, for example, after only five rows are inserted because the `cancel` is processed on the database before the stored procedure completes. SAP jConnect throws a `SQLException` when you try to close a `Statement` when there are still unprocessed results.

The missing inserts would not be reported to you. SAP jConnect programmers are strongly advised to adhere to these 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. In cases where you only executed `select` statements, this is fine. However, in cases where you executed `insert/update/delete` operations, this might 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.

## 2.11.4 Adjustments for Multithreading

Several threads simultaneously call methods on the same `Statement` instance, `CallableStatement`, or `PreparedStatement`, which SAP does not recommend you must manually synchronize the calls to the methods on the `Statement`; SAP `jConnect` 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 must synchronize the calls to the methods on the `Statement` or there might be conflicts.

## 2.11.5 `ResultSet.setCursorName`

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

Provided 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 documentation, all other SQL statements do not need to open a cursor and return a name.

For more information on how to use cursors in SAP `jConnect`, see *Cursors with Result Sets*.

## Related Information

[Use Cursors with Result Sets \[page 60\]](#)

## 2.11.6 Execute Stored Procedures

Executing 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.

Also, if you mix literals and question marks, you cannot use output parameters with a stored procedure.

This 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 this 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 these procedures:

- Use language commands, passing input parameters to them directly from Java variables using the `PreparedStatement` class.

```
// 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 SAP jConnect version 6.05 and later, use the `com.sybase.jdbcx.SybCallableStatement` interface:

```
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();
```

## 3 Security

SAP jConnect provides Secure Sockets Layer (SSL) and Kerberos options for securing client-server communications

- SSL – use SSL to encrypt communications, including the login exchange, between client and server applications.
- Kerberos – use Kerberos to authenticate Java applications or users of Java applications to SAP Adaptive Server Enterprise (SAP ASE) without sending user names or passwords over a network. Also use Kerberos to set up a single sign-on (SSO) environment and provide mutual authentication between the digital identity of a Java application and that of SAP ASE.

### i Note

You may use Kerberos to encrypt communications and provide data integrity checking, but these features have not been implemented for SAP jConnect.

You can Kerberos and SSL together, providing the advantage of both SSO and encryption of data transferred between client and server applications.

## Related Information

[Restrictions \[page 125\]](#)

[Implement Custom SSL Socket Plug-ins \[page 126\]](#)

[Kerberos \[page 130\]](#)

[Related Documents \[page 143\]](#)

[Securely Automate Database Administration Tasks Using Secure User Store \[page 143\]](#)

### 3.1 Restrictions

Kerberos and SSL is used with SAP Adaptive Server Enterprise (SAP ASE); SAP SQL Anywhere does not currently support either SSL or Kerberos security.

SAP recommends that you read related documentation about SSL and Kerberos before attempting to use either with SAP jConnect. The setup information assumes that the servers you intend to use have been configured to work properly with SSL, with Kerberos, or with both.

For further information on Kerberos, SSL, and configuring SAP ASE, see [Related Documents \[page 143\]](#). Also, see the white paper on setting up Kerberos, which is referenced in the *SAP jConnect for JDBC Release Bulletin*.



## 3.2 Implement Custom SSL Socket Plug-ins

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 can customize to enable encryption.

`com.sybase.jdbcx.SybSocketFactory` is a SAP jConnect extension interface that contains the `createSocket(String, int, Properties)` method, which returns a `java.net.Socket`. To use a custom socket factory in SAP jConnect, an application must implement this interface by defining the `createSocket()` method.

SAP jConnect uses the socket for subsequent input or output operations. Classes that implement `SybSocketFactory` create sockets and provide a general framework for the addition of public socket-level functionality, as shown:

```
/**
 * 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`, 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 can customize factories to return sockets with different networking timeouts 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`).

### i Note

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

## Related Information

[Using Custom Socket with jConnect \[page 127\]](#)

[Create and Configure a Custom Socket \[page 127\]](#)

[SSL Support in SAP jConnect \[page 129\]](#)

[TLS 1.1 and 1.2 Support in SAP jConnect \[page 130\]](#)

## 3.2.1 Using Custom Socket with jConnect

Review the steps to use custom socket with SAP jConnect.

### Procedure

1. Provide a Java class that implements `com.sybase.jdbcx.SybSocketFactory`.
2. Set the `SYB_SOCKET_FACTORY` connection property so that SAP jConnect can use your implementation to obtain a socket.

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

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

### Related Information

[Connection Properties \[page 14\]](#)

[Create and Configure a Custom Socket \[page 127\]](#)

## 3.2.2 Create and Configure a Custom Socket

You can create an instance of SSL socket and configure the socket, before SAP jConnect obtains it.

SAP jConnect uses the socket to connect to a server.

This example shows how an implementation of SSL can create an instance of `SSLSocket`, configure it, and then return it. 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 SAP jConnect

### Example

```
public class MySSLSocketFactory extends SSLSocketFactory
```

```

    implements SybSocketFactory
    {
    /**
    * Create a socket, set the cipher suites it can use, return
    * the socket.
    * Demonstrates how cipher 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);
Socket s =
    SSLSocketFactory.getDefault().createSocket
        (host, port);
// enable the cipher suites
((SSLSocket)s).setEnabledCipherSuites(enableThese);
// return the SSLSocket
return s;
}
// other methods
}

```

Because 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:

- `EncryptASE.java` – located in the `sample2` subdirectory of your SAP jConnect installation, this sample shows how to use the `SybSocketFactory` interface with SAP jConnect applications.
- `MySSLSocketFactoryASE.java` – also located in the `sample2` subdirectory of your SAP jConnect installation, this is a sample implementation of the `SybSocketFactory` interface that you can plug in to your application and use.

### 3.2.3 SSL Support in SAP jConnect

To use SSL sockets in versions of SAP jConnect earlier than 15.7 SP 100, you had to create an implementation of **SybSocketFactory** interface and use it by setting the `SYB_SOCKET_FACTORY` connection property.

Starting from version 15.7 SP 100, SAP jConnect has built-in support to connect to SAP Adaptive Server Enterprise (SAP ASE) using SSL sockets. The new connection property `ENABLE_SSL` when set to:

- `false` – (the default) SAP jConnect will not use SSL sockets.
- `true` – SAP jConnect uses SSL sockets and the target SAP ASE must be enabled for SSL socket connections. SAP jConnect ignores the `SYB_SOCKET_FACTORY` connection property.

#### i Note

SAP recommends that you set the login timeout using `DriverManager.setLoginTimeout` property to allow the connection to timeout when attempting SSL connection on a non SSL enabled SAP ASE.

The SSL socket feature depends on the following standard Java properties:

- `javax.net.ssl.keyStore`
- `javax.net.ssl.keyStorePassword`

- `javax.net.ssl.trustStore`
- `javax.net.ssl.trustStorePassword`
- `javax.net.ssl.trustStore`
- `javax.net.ssl.trustStoreType`

See the Java J2SE 6 Documentation for more information on Java standard properties.

### 3.2.4 TLS 1.1 and 1.2 Support in SAP jConnect

A new connection property, `SSL_MIN_PROTOCOL_VERSION`, specifies the minimum SSL protocol version to establish a connection.

The `SSL_MIN_PROTOCOL_VERSION` property, when set, requests the JCE provider to set the TLS protocol to establish a handshake with the SSL-enabled server. During the handshake between JCE and the SSL server, the protocol with the highest version is given priority and the connection between the client and the server is established. The connection fails if the protocol handshake does not happen with the minimum specified protocol version.

The valid protocol versions, from highest to lowest priority, are:

- TLSv1.2
- TLSv1.1
- TLSv1
- SSLv3
- SSLv2Hello

To use the `SSL_MIN_PROTOCOL_VERSION` property, set `ENABLE_SSL` to `TRUE`.

The default value of `SSL_MIN_PROTOCOL_VERSION` is `TLSv1`.

#### ❖ Example

```
SSL_MIN_PROTOCOL_VERSION = TLSv1.1
```

In this case, the JCE provider first attempts a handshake with the protocol version TLSv1.2 (highest). When this fails, it tries using TLSv1.1. If the handshake with TLSv1.1 also fails, then the provider stops trying to establish a connection. The JCE does not attempt the handshake with the other protocols of lower priorities like TLSv1, SSLv3, and SSLv2Hello, even if they are supported by the JRE.

## 3.3 Kerberos

Kerberos is a network authentication protocol that uses encryption for authentication of client-server applications.

Kerberos provides these advantages for users and system administrators:

- A Kerberos database can serve as a centralized storehouse for users.

- Kerberos facilitates the single-sign-on (SSO) environment, in which a user system login provides the credentials necessary to access a database.
- Kerberos is an IETF standard. Interoperability is possible between different implementations of Kerberos.

## Related Information

[Configuring Kerberos for SAP jConnect \[page 131\]](#)

[GSSMANAGER\\_CLASS Connection Property \[page 132\]](#)

[Kerberos Environment \[page 135\]](#)

[Sample Applications \[page 139\]](#)

[Interoperability \[page 141\]](#)

[Troubleshooting Kerberos \[page 142\]](#)

### 3.3.1 Configuring Kerberos for SAP jConnect

Review the instructions to configure SAP jConnect to use Kerberos security mechanism.

## Prerequisites

There are several prerequisites for configuring Kerberos for SAP jConnect:

- JDK 6 or later
- A Java Generic Security Services (GSS) Manager:
  - The default GSS Manager, which is part of the JDK, or
  - Wedgetail JCSI Kerberos version 2.6 or later, or
  - CyberSafe TrustBroker Application Security Runtime Library version 3.1.0 or later, or
  - A GSS Manager implementation from another vendor.
- A key distribution center(KDC) that is supported and interoperable at the server side with your GSS library and at the client side with your GSSManager.

## Procedure

1. Set the `REQUEST_KERBEROS_SESSION` property to true.
2. Set the `SERVICE_PRINCIPAL_NAME` property to the name that your SAP Adaptive Server Enterprise is running under. In general, this is the name set with the `-s` option when the server is started. The service principal name must also be registered with the KDC. If you do not set a value for this property, SAP jConnect uses the host name of the client machine.
3. (Optional) Set the `GSSMANAGER_CLASS` property.

For more information on the `REQUEST_KERBEROS_SESSION` and `SERVICE_PRINCIPAL_NAME`, see the [jConnect Connection Properties \[page 15\]](#)

## Related Information

[GSSMANAGER\\_CLASS Connection Property \[page 132\]](#)  
[Programming Information \[page 8\]](#)

### 3.3.2 GSSMANAGER\_CLASS Connection Property

When using Kerberos, SAP jConnect relies on several Java classes that implement the Generic Security Services (GSS) API.

Much of this functionality is provided by the `org.ietf.jgss.GSSManager` class.

SAP jConnect checks the value of `GSSMANAGER_CLASS` for a `GSSManager` class object to use in Kerberos authentication.

If the value of `GSSMANAGER_CLASS` is set to a string instead of a class object, SAP jConnect uses the string to create an instance of the specified class and uses the new instance in Kerberos authentication.

If the value of `GSSMANAGER_CLASS` is set to something that is neither a `GSSManager` class object nor a string, or if SAP jConnect encounters a `ClassCastException`, jConnect throws a `SQLException` indicating the problem.

Java allows vendors to provide their own implementations of the `GSSManager` class.

Examples of vendor-supplied `GSSManager` implementations are those provided by Wedgetail Communications and CyberSafe Limited. Users can configure a vendor-written `GSSManager` class to work in a particular Kerberos environment. Vendor-supplied `GSSManager` classes may also offer more interoperability with Windows than the standard Java `GSSManager` class provides.

Before using a vendor-supplied implementation of `GSSManager`, be sure to read the vendor documentation. Vendors use system property settings other than the standard Java system properties used for Kerberos and may locate realm names and key distribution center (KDC) entries without using configuration files.

## Related Information

[Setting Up the GSSMANAGER\\_CLASS Property \[page 133\]](#)

### 3.3.2.1 Setting Up the GSSMANAGER\_CLASS Property

Use a vendor implementation of `GSSManager` with SAP jConnect by setting the `GSSMANAGER_CLASS` connection property.

#### Context

There are two ways to set this property:

- Create an instance of `GSSManager`, and set this instance as the value of the `GSSMANAGER_CLASS` property.
- Set the value of the `GSSMANAGER_CLASS` property as a string, specifying the fully qualified class name of the `GSSManager` object. SAP jConnect uses this string to call `Class.forName().newInstance()` and casts the returned object as a `GSSManager` class.

In either case, the application `CLASSPATH` variable must include the location of the classes and `.jar` files for the vendor implementation.

#### Note

If you do not set the `GSSMANAGER_CLASS` connection property, SAP jConnect uses the `org.ietf.jgss.GSSManager.getInstance` method to load the default Java `GSSManager` implementation.

#### Example

When you use the `GSSMANAGER_CLASS` connection property to pass in a fully qualified class name, SAP jConnect calls the no-argument constructor for the `GSSManager`. This instantiates a `GSSManager` that is in the default configuration for the vendor implementation, so you do not have control over the exact configuration of the `GSSManager` object. If you create your own instance of `GSSManager`, you can use constructor arguments to set configuration options.

#### Related Information

[GSS Manager Examples \[page 134\]](#)



### 3.3.2.1.1 GSS Manager Examples

Review instructions to create an instance of `GSSManager` for your requirement or allow SAP jConnect to create a `GSSManager` object when the `GSSMANAGER_CLASS` connection property is set to a fully qualified class name.

#### Related Information

[Creating an Instance of GSSManager \[page 134\]](#)

[Passing String to GSSMANAGER\\_CLASS \[page 135\]](#)

#### 3.3.2.1.1.1 Creating an Instance of GSSManager

Create an instance of `GSSManager` and pass it to the `GSSMANAGER_CLASS` property.

#### Procedure

1. Instantiate a `GSSManager` in your application code:

```
GSSManager gssMan = new com.dstc.security.kerberos.gssapi.GSSManager();
```

This example uses the default constructor with no arguments. You can use other vendor-supplied constructors, which allow you to set various configuration options.

2. Pass the new `GSSManager` instance into the `GSSMANAGER_CLASS` connection property:

```
Properties props = new Properties();  
props.put("GSSMANAGER_CLASS", gssMan);
```

3. Use these connection properties, including `GSSMANAGER_CLASS`, in your connection:

```
Connection conn = DriverManager.getConnection (url, props);
```

### 3.3.2.1.1.2 Passing String to GSSMANAGER\_CLASS

Pass string to GSSMANAGER\_CLASS in an application.

#### Procedure

1. Create a string specifying the fully qualified class name of the GSSManager object. For example:

```
String gssManClass = "com.dstc.security.kerberos.gssapi.GSSManager";
```

2. Pass the string to the GSSMANAGER\_CLASS connection property. For example:

```
Properties props = new Properties();  
props.put("GSSMANAGER_CLASS", gssManClass);
```

3. Use these connection properties, including GSSMANAGER\_CLASS, in your connection. For example:

```
Connection conn = DriverManager.getConnection (url, props);
```

### 3.3.3 Kerberos Environment

You can use SAP jConnect with three different implementations of Kerberos.

- CyberSafe
- MIT
- Microsoft Active Directory

See the [Kerberos white paper](#) .

#### Related Information

[CyberSafe \[page 136\]](#)

[MIT \[page 137\]](#)

[Microsoft Active Directory \[page 137\]](#)

### 3.3.3.1 CyberSafe

Review the CyberSafe Kerberos implementation in SAP jConnect.

#### Encryption keys

specify a Data Encryption Standard (DES) key when creating a principal to be used by Java in the CyberSafe KDC.

The Java reference implementation does not support Triple Data Encryption Standard (3DES) keys.

#### i Note

You can use 3DES keys if you are using CyberSafe GSSManager with a CyberSafe KDC and have set the `GSSMANAGER_CLASS` property.

#### Address mapping and realm information

CyberSafe uses DNS records to locate KDC address mapping and realm information.

CyberSafe Kerberos does not use a `krb5.conf` configuration file. Alternately, CyberSafe locates KDC address mapping and realm information in the `krb.conf` and `krb.realms` files, respectively. See CyberSafe documentation for more information.

If you are using the standard Java `GSSManager` implementation, you must still create a `krb5.conf` file for use by Java. The CyberSafe `krb.conf` file is formatted differently from the `krb5.conf` file. Create a `krb5.conf` file as specified in the Java SE documentation or in the MIT documentation. You do not need a `krb5.conf` file if using the CyberSafe `GSSManager`.

For examples of the `krb5.conf` file, see white paper on setting up Kerberos, the URL is referenced in the *SAP jConnect for JDBC Release Bulletin*.

#### Solaris

when using CyberSafe client libraries on Solaris, make sure your library search path includes the CyberSafe libraries before any other Kerberos libraries.

#### ❖ Example

A client uses `krb5.conf` file with a CyberSafe or MIT KDC. For example:

```
# Please note that customers must alter the
# default_realm, [realms] and [domain_realm]
# information to reflect their Kerberos environment.
# Customers should *not* attempt to use this file as is.
#
[libdefaults]
    default_realm = ASE
    default_tgs_enctypes = des-cbc-crc
    default_tkt_enctypes = des-cbc-crc
    kdc_req_checksum_type = 2
    ccache_type = 2
[realms]
    ASE = {
        kdc = kdchost
        admin_server = kdchost
    }
[domain_realm]
    .sybase.com = ASE
    sybase.com = ASE
[logging]
```

```

        default = FILE:/var/krb5/kdc.log
        kdc = FILE:/var/krb5/kdc.log
    kdc_rotate = {
    # How often to rotate kdc.log. Logs will get rotated
    # no more often than the period, and less often if the
    # KDC is not used frequently.
        period = 1d
    # how many versions of kdc.log to keep around
    # (kdc.log.0, kdc.log.1, ...)
        versions = 10
    }
[appdefaults]
    kinit = {
        renewable = true
        forwardable = true
    }

```

### 3.3.3.2 MIT

Specify a DES key when creating a principal to be used by Java in the MIT KDC.

The Java reference implementation does not support 3DES keys.

If you plan to use only the standard Java `GSSManager` implementation, specify an encryption key of type `des-cbc-crc` or `des-cbc-md5`. Specify the encryption type as:

```
des-cbc-crc:normal
```

where `normal` is the type of key salt. It may be possible to use other salt types.

#### i Note

If you are using `Wedgetail GSSManager`, you can create principals in an MIT KDC of type `des3-cbc-sha1-kd`.

### 3.3.3.3 Microsoft Active Directory

Review the components in a Microsoft Active Directory server for Kerberos.

#### User accounts and service principal

make sure that you have set up accounts in Active Directory for your user principals user(the users) and service principals (the accounts that represent your database servers). Your user principals and service principals should both be created as Users within Active Directory.

#### Client machines


modify the Windows Registry to use the Java reference implementation to set up an SSO environment.

See the instructions at the [Microsoft support site](#) ➡ to modify Windows Registry.

**Configuration file** on Windows, the Kerberos configuration file is called `krb5.ini`. Java looks for `krb5.ini` by default at `C:\WINNT\krb5.ini`.

Java allows you to specify the location of this file. The format of `krb5.ini` is identical to that of `krb5.conf`.

For examples of the `krb5.conf` file, see Kerberos white paper, which is referenced in the *SAP jConnect for JDBC Release Bulletin*.

For more information on Kerberos for Microsoft Active Directory, see the [Microsoft Developer Network](#) .

## ❁ Example

A client uses the `krb5.conf` file with Active Directory as the KDC. For example:

```
# Please note that customers must alter the
# default_realm, [realms] and [domain_realm]
# information to reflect their Kerberos environment.
# Customers should *not* attempt to use this file as is.
#
[libdefaults]
    default_realm = W2K.SYBASE.COM
    default_tgs_enctypes = des-cbc-crc
    default_tkt_enctypes = des-cbc-crc
    kdc_req_checksum_type = 2
    ccache_type = 2
[realms]
    W2K.SYBASE.COM = {
        kdc = 1.2.3.4:88
        admin_server = adserver
    }
[domain_realm]
    .sybase.com = W2K.SYBASE.COM
    sybase.com = W2K.SYBASE.COM
[logging]
    default = FILE:/var/krb5/kdc.log
    kdc = FILE:/var/krb5/kdc.log
    kdc_rotate = {
# How often to rotate kdc.log. Logs will get rotated no
# more often than the period, and less often if the KDC
# is not used frequently.
        period = 1d
# how many versions of kdc.log to keep around
# (kdc.log.0, kdc.log.1, ...)
        versions = 10
    }
[appdefaults]
    kinit = {
        renewable = true
        forwardable = true
    }
```

## Related Information

[Setting DES Encryption \[page 139\]](#)

### 3.3.3.3.1 Setting DES Encryption

If you intend to use the Java reference GSS Manager implementation, you must use DES encryption for both user and service principals.

#### Procedure

1. In the Active Directory, right-click on the specific user principal or service principal name.
2. Select **Properties**.
3. Click the **Account** tab.
4. For both the user principal and service principal, specify that DES encryption types should be used.

### 3.3.4 Sample Applications

The two commented code samples available in the `jConnect-16_0/sample2` directory illustrate how to establish a Kerberos connection to SAP Adaptive Server Enterprise (SAP ASE).

- `ConnectKerberos.java` – a simple Kerberos login to SAP ASE.
- `ConnectKerberosJAAS.java` – a more detailed sample showing how a Kerberos login might be implemented within application-server code.

#### Related Information

[Running ConnectKerberos.java \[page 139\]](#)

### 3.3.4.1 Running ConnectKerberos.java

Review the instructions to run `ConnectKerberos.java` file sample application.

#### Procedure

1. Make sure your machine has valid Kerberos credentials. This task varies depending on your machine and environment.
  - Windows – you can establish Kerberos credentials for a machine in an Active Directory environment by successfully logging in using Kerberos authentication.

- UNIX or Linux – you can establish Kerberos credentials for a UNIX or Linux machine using the `kinit` utility for your Kerberos client. If you do not obtain an initial credential using `kinit`, you are prompted for a user name and password when you attempt to run the sample application.

### **i Note**

Typically, the GSSManager provider implementation provided by standard JDK can use only the DES\_CBC\_MD5 and DES\_CBC\_CRC encryption types. You may be able to use other encryption types by using third-party software and setting the `GSSMANAGER_CLASS` property.

2. Determine the location of the credentials for your machine.
  - Windows – for a machine running in an Active Directory environment, Kerberos credentials are stored in an in-memory ticket cache.
  - UNIX or Linux – for a UNIX or Linux machine using the JRE supplied, CyberSafe, Solaris, or MIT implementations of Kerberos, `kinit` places credentials by default in `/tmp/krb5cc_{user_id_number}`, where `<{user_id_number}>` is unique to your user name.

If the credentials are placed elsewhere, specify that location in the `sample2/exampleLogin.conf` file by setting the `ticketCache` property.

3. Specify to the Java reference implementation the default realm and host name of the KDC machine. Java may obtain this information from the `krb5.conf` or `krb5.ini` configuration files or from Java System properties. If you use a vendor GSS Manager implementation, that implementation may obtain host and realm information from DNS SRV records.

SAP recommends that you use a Kerberos configuration file, which allows for more control of the Kerberos environment, including the ability to specify to Java the type of encryption to request during authentication.

### **i Note**

On Linux, the Java reference implementation looks for the Kerberos configuration file in `/etc/krb5.conf`.

If you do not use a Kerberos configuration file, and your Kerberos configuration is not set up to use DNS SRV records, you can specify the realm and KDC using the `java.security.krb5.realm` and `java.security.krb5.kdc` system properties.

4. Edit `ConnectKerberos.java` so that the connection URL points to your database.
5. Compile `ConnectKerberos.java`.

Ensure that you are using JDK version 6 or later. Read through the source code comments, and ensure the `jconn4.jar` from your `jConnect` installation is specified in your `CLASSPATH` environment variable.

6. Execute `ConnectKerberos.class`:

```
java ConnectKerberos
```

Ensure that you are using the Java version 6 executable. The sample application output explains that a successful connection has been established and executes the SQL:

```
select 1
```

- To execute the sample without using a Kerberos configuration file, use:

```
java -Djava.security.krb5.realm=<your_realm>
-Djava.security.krb5.kdc=<your_kdc> ConnectKerberos
```

where <your\_realm> is your default realm, and <your\_kdc> is your KDC.

- If necessary, you can run the sample application in debug mode to see debug output from the Java Kerberos layer:

```
java -Dsun.security.krb5.debug=true ConnectKerberos
```

You can also make a Kerberos connection using `IsqlApp`, the Java version of `isql`, located in the `jConnect-16_0/classes` directory:

```
java IsqlApp -S jdbc:sybase:Tds:<hostName>:<portNum>
-K <service_principal_name>
-F <path_to_JAAS_login_module_config_file>
```

## 3.3.5 Interoperability

SAP jConnect supports interoperability combinations of KDCs, GSS libraries, and platforms on which SAP jConnect has successfully established a connection to SAP Adaptive Server Enterprise (SAP ASE).

The absence of any particular combination does not indicate that a connection cannot be established with that combination. You can find the most recent status at the [jConnect for JDBC Web site](#).

Interoperability Combinations

Client Platform	KDC	GSSManager	GSS C libraries <sup>a</sup>	SAP ASE Platform
Solaris 8 <sup>b</sup>	CyberSafe	Java GSS	CyberSafe	Solaris 8
Solaris 8	Active Directory <sup>c</sup>	Java GSS	CyberSafe	Solaris 8
Solaris 8	MIT	Java GSS	CyberSafe	Solaris 8
Solaris 8	MIT	Wedgetail GSS <sup>d</sup>	MIT	Solaris 8
Solaris 8	CyberSafe	Wedgetail GSS <sup>e</sup>	CyberSafe	Solaris 8
Windows 2000	Active Directory	Java GSS	CyberSafe	Solaris 8
Windows XP	Active Directory	Java GSS <sup>f</sup>	CyberSafe	Solaris 8



Client Platform	KDC	GSSManager	GSS C libraries \ <sup>a</sup>	SAP ASE Platform
<p>a. These are the libraries that SAP ASE is using for its GSS functionality.</p> <p>b. All Solaris 8 platforms in this table are 32-bit.</p> <p>c. All Active Directory entries in the table refer to an Active Directory server running on Windows 2000. For Kerberos interoperability, Active Directory users must be set to "Use DES encryption types for this account."</p> <p>d. Used Wedgetail JCSI Kerberos 2.6. The encryption type was 3DES.</p> <p>e. Used Wedgetail JCSI Kerberos 2.6. The encryption type was DES.</p> <p>f. Java 1.4.x has a bug which requires that clients use <code>System.setProperty("os.name", "Windows 2000");</code> to ensure that Java can find the in-memory credential on Windows XP clients.</p>				

SAP recommends that you use the latest versions of these libraries. Contact the vendor if you intend to use older versions or if you have problems with non-SAP products.

## Related Information

[Encryption Types \[page 142\]](#)

### 3.3.5.1 Encryption Types

The standard Java GSS implementation provided by typical JREs supports only DES encryption.

If you intend to use the 3DES, RC4-HMAC, AES-256, or AES-128 encryption standards, you must use the CyberSafe or Wedgetail GSSManagers.

Refer to the respective documentation for more information about Wedgetail and CyberSafe.

### 3.3.6 Troubleshooting Kerberos

Review the considerations when troubleshooting Kerberos security issues.

- The Java reference implementation supports only the DES encryption type. You must configure your Active Directory and KDC principals to use DES encryption.
- The value of the `SERVICE_PRINCIPAL_NAME` property must be set to the same name you specify with the `-s` option when you start your data server.
- Check the `krb5.conf` and `krb5.ini` files. For CyberSafe clients, check the `krb.conf` and `krb.realms` files or DNS SRV records.
- You can set the `debug` property to `true` in the JAAS login configuration file.
- You can set the `debug` property to `true` at the command line:

```
-Dsun.security.krb5.debug=true
```

- The JAAS login configuration file provides several options that you can set for your particular needs. For information about JAAS and the Java GSS API, refer to:
  - *JAAS Login Configuration File*
  - *Class Krb5LoginModule*
  - *Troubleshooting JGSS*

## 3.4 Related Documents

Review the additional information on Kerberos security.

- *Java tutorial on JAAS and the Java GSS API*
- *MIT Kerberos documentation and download site*
- *CyberSafe Limited*
- *CyberSafe Limited document on Windows-Kerberos interoperability*
- *Kerberos RFC 1510*

## 3.5 Securely Automate Database Administration Tasks Using Secure User Store

With secure user store feature, you can connect to the SAP ASE server without explicitly entering username and password in the command line interface; instead, you can establish the connection between a client application and the SAP ASE server using a user-generated key that retrieves the connection information (username, password, host, and port number) from the secure store.

Use the `aseuserstore` utility to create and manage the secure store, generate user keys, and add those keys to the secure store. The feature is supported for ADO.NET, ODBC, jConnect, and SAP ASE native client applications.

Each client category (for example, ADO.NET or ODBC) requires its own entry in the secure store, even when the connection parameters are the same.

This password-less connection option lets you automate database administration tasks without compromising data confidentiality. The `aseuserstore` utility is located in:

- (UNIX) `$SYBASE/OCS-16_0/bin`
- (Windows) `%SYBASE%\OCS-16_0\bin`

## Related Information

[How Does the aseuserstore Utility Work? \[page 144\]](#)

[Enabling Secure User Store for jConnect Client \[page 144\]](#)

## 3.5.1 How Does the aseuserstore Utility Work?

The `aseuserstore` utility provides commands to store client connection information in the secure store.

Use the `aseuserstore` utility commands to insert, list, and delete connection information in the form of user keys in the secure store. Each user key consists of a username and password, and optionally a hostname, port number, and a database name.

Users are encouraged to periodically change the secure store's encryption key using the `aseuserstore` utility's `CHANGEKEY` command. If there are changes to the client connection parameters (change in the port number, for instance), re-run the `aseuserstore` utility for that key with its new key value (that is, the new port number). Its corresponding entry is then updated in the secure store.

### i Note

The `aseuserstore` utility stores the secure user store in `$HOME/.ase/<hostname>`, where `<hostname>` is the name of the host where you have executed the command. If you need to override the default hostname, use the utility's `-H` option.

## 3.5.2 Enabling Secure User Store for jConnect Client

A new connection property, `SECURE_STORE_KEY`, is added to the jConnect JDBC driver to specify the keyname that retrieves the user credentials from the secure store.

### i Note

When using the secure store option, username and password are always retrieved from the secure store. If you specify the username and/or password in the connection string also, it results in an error. Other connection parameters are optional, and may be specified in the connection string if they are not set in the secure store. If you specify connection parameters in both the connection string and secure store, their values must match for the connection to succeed.

### Example 1

The following example shows the usage of `Properties` object to specify the `SECURE_STORE_KEY` property.

### ≡ Code Syntax

```
Properties props = new Properties();
props.put("SECURE_STORE_KEY", "DURIANDB");
Connection conn =
    DriverManager.getConnection("jdbc:sybase:Tds:durian.sjc.sap.corp:1600", props);
```

### Example 2

The following example shows the usage of `SECURE_STORE_KEY` property through a connection URL string. Here, the connection will be established with `localhost:1600` if the connection information for `ASE1TESTUSER` is also specified as `localhost:1600`, else it results in an error because there is a mismatch between secure store and the connection string.

#### Code Syntax

```
DriverManager.getConnection("jdbc:sybase:Tds:localhost:1600?  
SECURE_STORE_KEY=ASE1TESTUSER);
```

#### Example 3

#### Code Syntax

In the following example, the host and port information is not provided in the connection string. In this scenario, the driver uses the connection information present in the secure store.

```
DriverManager.getConnection("jdbc:sybase:Tds:?SECURE_STORE_KEY=ASE1TESTUSER);
```

## 3.5.3 Examples of Creating User Keys for jConnect Clients

Example 1: Stores a user key (`mykey`) for the database user (`user1`) with the password (`password`).

```
aseuserstore set mykey myhost:1600 user1 password
```

Note that the connection information can be left blank by specifying empty quotes. In this scenario, the connection information provided in the connection string is used.

```
aseuserstore set mykey "" user1 password
```

Example 2: Stores a user key (`mykey`) for the database user (`user1`) without specifying the password. The user is then prompted to enter the password. Note that the entered password is not displayed on the terminal.

```
aseuserstore set mykey 'myhost:1600' user1  
Password:
```

Example 3: Stores a user key (`mykey`) for the database user (`user1`). The command, in addition to the username and password, also includes the database (`database1`), which is specified using the `/` switch.

```
aseuserstore set mykey 'myhost:1600/database1' user1 password
```

Example 4: Stores a user key (`mykey`) for the database user (`user1`). The command, in addition to the username and password, also includes the database (`database1`), which is specified using the `@` switch.

```
aseuserstore set mykey myhost:1600@database1 user1 password
```

# 4 Troubleshooting

Review the solutions and workarounds for problems you might encounter when using SAP jConnect.

## Related Information

[Debugging with SAP jConnect \[page 146\]](#)  
[Dynamic Logging \[page 150\]](#)  
[Capture TDS Communication \[page 152\]](#)  
[Resolve Connection Errors \[page 154\]](#)  
[Manage Memory in SAP jConnect Applications \[page 155\]](#)  
[Resolve Stored Procedure Errors \[page 156\]](#)  
[Resolve Custom Socket Implementation Error \[page 157\]](#)

## 4.1 Debugging with SAP jConnect

SAP jConnect includes a `Debug` class that contains a set of debugging functions.

### Context

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 SAP jConnect installation also includes a complete set of debug-enabled classes. These classes are located in the `devclasses` subdirectory under your SAP jConnect installation directory. For debugging purposes, you must redirect your `CLASSPATH` environment variable to reference the debug mode runtime classes (`devclasses/jconn4d.jar`), rather than the standard SAP jConnect `classes` directory. You can also do this by explicitly providing a `-classpath` argument to the `java` command when you run a Java program.

## Related Information

[Obtaining an Instance of the Debug Class \[page 147\]](#)  
[Turning On Debugging in an Application \[page 147\]](#)  
[Turning Off Debugging in an Application \[page 148\]](#)

[Setting the CLASSPATH for Debugging \[page 148\]](#)

[Using the Debugging Methods \[page 149\]](#)

## 4.1.1 Obtaining an Instance of the Debug Class

Import the `Debug` interface and obtain an instance of the `Debug` class by calling the `getDebug` method on the `SybDriver` class.

### Context

```
import com.sybase.jdbcx.Debug;
//
...
SybDriver sybDriver = (SybDriver)
Class.forName("com.sybase.jdbc4.jdbc.SybDriver").newInstance();
Debug sybdebug = sybDriver.getDebug();
...
```

## 4.1.2 Turning On Debugging in an Application

Use the `debug` method on the `Debug` object to turn on debugging within your application.

### Context

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")
```

Using "STATIC" in the class string turns on debugging for all static methods in SAP 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`.

## 4.1.3 Turning Off Debugging in an Application

Review the instruction to turn off debugging method.

### Context

Add this call:

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

## 4.1.4 Setting the CLASSPATH for Debugging

Before you run your debug-enabled application, replace the optimized SAP jConnect `jconn4.jar` file with the debug version `jconn4d.jar`, which you can find in the `devclasses` subdirectory under your SAP jConnect installation directory.

### Context

To set the environment variable:

- For UNIX, replace `$JDBC_HOME/classes/jconn4.jar` with `$JDBC_HOME/devclasses/jconn4d.jar`.
- For Windows, replace `%JDBC_HOME%\classes\jconn4.jar` with `%JDBC_HOME%\devclasses\jconn4d.jar`.

## 4.1.5 Using the Debugging Methods

Customize the debugging methods in SAP jConnect.

### Context

You can add calls to other `Debug` methods.

#### ❖ Example

If any of these methods are static, use null for the object parameter.

- `println` – 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` – 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` and `stopTimer` – 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
```

## 4.2 Dynamic Logging

Starting with 15.7 ESD #4, SAP jConnect for JDBC supports logging mechanism by implementing the standard Java logger mechanism.

### Example

The application obtains the handle of the SAP jConnect logger, and turn logging on or off as and when required.

```
try  
{  
    // Get logger for all classes present in  
    // "com.sybase.jdbc4.jdbc" package  
    Logger LOG = Logger.getLogger("com.sybase.jdbc4.jdbc");  
    // To log class-specific log message,  
    // provide complete class name, for example:  
    //Logger.getLogger("com.sybase.jdbc4.jdbc.  
    //SybConnection");  
    //Get handle as per user's requirement  
    Handler handler = new ConsoleHandler();  
    //Set logging level  
    handler.setLevel(Level.ALL);  
    //Added user specific handler to logger object  
    LOG.addHandler(handler);  
    //Set logging level  
    LOG.setLevel(Level.ALL);  
    Class.forName("com.sybase.jdbc4.jdbc.SybDriver");  
    Properties properties = new Properties();  
    properties.put("USER", USER_NAME);  
    properties.put("PASSWORD", PASSWORD);  
    Connection con = DriverManager.getConnection("jdbc:sybase:Tds:" +  
        HOST_PORT, properties);  
    Statement stmt = con.createStatement();  
    stmt.execute("select @@version");  
    //Dynamically turn off logging mechanism  
    LOG.setLevel(Level.OFF);  
    con.close();  
    ...  
}
```

```
}
```

## Logging Levels

SAP jConnect allows application users to set message granularity to Level.FINE, Level.FINER, and Level.FINEST. For example:

- When a user sets the logging level to Level.FINE on SybConnection class, jConnect reports: Dr1\_Col setClientInfo(Properties)
- Level.FINER on SybConnection class reports: Dr1\_Col setClientInfo(Properties.size = [3])
- Level.FINEST on SybConnection class reports: Dr1\_Col setClientInfo(Properties = [[ClientUserValue, ApplicationNameValue, ClientHostnameValue]])

## Related Information

[Enabling Logging Dynamically in SAP jConnect \[page 151\]](#)

[Enabling Logging Statically in SAP jConnect \[page 152\]](#)

## 4.2.1 Enabling Logging Dynamically in SAP jConnect

Dynamically enable logging by using LogHandler API programmatically in your application.

### Procedure

1. Enable or disable logging by programmatically using LogHandler API. Enter the following code to enable console-level logging on SybConnection and SybStatement classes:

```
LogManager logManager = LogManager.getLogManager();
Handler handler = new ConsoleHandler();
handler.setLevel(Level.ALL);
Logger connLOG = Logger.getLogger(SybConnection.class.getName());
connLOG.addHandler(handler);
connLOG.setLevel(Level.FINE);
logManager.addLogger(connLOG);
Logger stmtLOG = Logger.getLogger(SybStatement.class.getName());
stmtLOG.addHandler(handler);
stmtLOG.setLevel(Level.FINE);
logManager.addLogger(stmtLOG);
```

2. Run the application.
3. You can tune the logging setting dynamically using the LogHandler as given in Step 1. Some of the suggested ways to achieve this:

- Expose an interface where you can change a property that internally manages logging levels.
- Run a surrogate thread that tunes the logging as required in your application.

## 4.2.2 Enabling Logging Statically in SAP jConnect

Statically enable logging in SAP jConnect that implements the standard Java logging mechanism.

### Procedure

1. Use the text editor to modify the contents of the standard logging file in `$JRE_DIR/lib/logging.properties`.

```
handlers= java.util.logging.FileHandler
java.util.logging.FileHandler.formatter = com.sybase.jdbc4.utils.LogUtil
.level= ALL
```

2. Add or enter the following contents in the file:

```
com.sybase.jdbc4.jdbc.SybDriver.level = FINEST
com.sybase.jdbc4.jdbc.SybConnection.level = FINEST
com.sybase.jdbc4.jdbc.SybStatement.level = FINER
com.sybase.jdbc4.jdbc.SybPreparedStatement.level = FINE
com.sybase.jdbc4.jdbc.SybResultSet.level = FINE
```

3. Adjust the level of logging to Level.FINE, Level.FINER, and Level.FINEST to set the appropriate granularity of logging.

#### **i** Note

SAP jConnect does not support package-level logging.

4. Save the `logging.properties` file.

## 4.3 Capture TDS Communication

TDS is the SAP jConnect-proprietary protocol for handling communication between a client application and SAP Adaptive Server Enterprise (SAP ASE).

SAP 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, 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 SAP Technical Support for analysis.

### Note

The captured TDS protocol data saved to a file contains sensitive user authentication information and may contain confidential company or customer data. To protect this confidential data from unauthorized or accidental disclosure, use file permissions or encryption to properly protect the files containing captured data.

## Related Information

[PROTOCOL\\_CAPTURE Connection Property \[page 153\]](#)

[Pause and Resume Methods in Capture Class \[page 153\]](#)

### 4.3.1 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 SAP Adaptive Server Enterprise (SAP ASE).

`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.

This example shows the use of `PROTOCOL_CAPTURE` to send TDS data to the file `tds_data`:

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

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

### 4.3.2 Pause and Resume Methods in Capture Class

The `Capture` class is in the `com.sybase.jdbcx` package, and contains `pause` and `resume` methods.

`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. You can limit the size of the capture file, if you know where in an application you want to capture TDS data.

## Related Information

[Limiting Size of Capture File \[page 154\]](#)

### 4.3.2.1 Limiting Size of Capture File

Review the instructions to limit the size of the capture file.

#### Procedure

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` where you want to start capturing TDS data.
3. Place `cap.pause` where you want to stop capturing data.

## 4.4 Resolve Connection Errors

Address the problems that may arise when you are trying to establish a connection or start a gateway.

```
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 SAP Adaptive Server Enterprise (SAP ASE). 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.”

## 4.5 Manage Memory in SAP jConnect Applications

Use the `Statement` objects and subclasses, if you notice increased memory use in SAP jConnect applications

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

For example, this statement causes problems:

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

Instead, use:

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

- Native support for scrollable or updatable scrollable cursors may not be available, depending on the version of SAP Adaptive Server Enterprise (SAP ASE) or SAP SQL Anywhere database you are connecting to. To support scrollable or updatable scrollable cursors when not supported natively by the back-end server, SAP 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 client memory. Because this may cause a performance degradation, SAP recommends that you use `TYPE_SCROLL_INSENSITIVE` result sets only when the result set is reasonably small. SAP jConnect determines if the SAP ASE connection supports native scrollable cursor functionality and uses it instead of client-side caching. As a result, most applications can expect significant performance gain in accessing out-of-order rows and reduction in client-side memory requirements.

## 4.6 Resolve Stored Procedure Errors

Address the problems that occur when you are trying to use SAP jConnect and stored procedures.

### Related Information

[RPC Returns Fewer Output Parameters Than Registered \[page 156\]](#)

[Fetch/State Error \[page 156\]](#)

[Stored Procedure Executed in Unchained Transaction Mode \[page 157\]](#)

### 4.6.1 RPC Returns Fewer Output Parameters Than Registered

If you call `CallableStatement.registerOutParam` for more parameters than you have declared as OUTPUT parameters in the stored procedure, an error occurs.

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

Make sure you have declared all of the appropriate parameters as "OUTPUT." Look at the line of code that reads:

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

#### **i** Note

If you receive this error while using SAP SQL Anywhere, upgrade to SAP SQL Anywhere version 5.5.04 or later.

### 4.6.2 Fetch/State Error

Fetch/State error occurs if a query does not return row data.

You can use the `CallableStatement.executeUpdate` or `execute` methods rather than the `executeQuery` method.

As required by the JDBC standards, SAP jConnect throws a SQL exception if `executeQuery` has no result sets.

### 4.6.3 Stored Procedure Executed in Unchained Transaction Mode

This error occurs when JDBC attempts to send the connection in `autocommit(true)` mode.

SAP Adaptive Server Error 7713 - Stored Procedure can only be executed in unchained transaction 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_procxmode <procedure_name,>"anymode"
```

## 4.7 Resolve Custom Socket Implementation Error

Custom socket implementation error occurs when you try 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.



# 5 Performance and Tuning

Review the instructions to fine-tune or improve performance when working with SAP jConnect.

## Related Information

[Improve SAP jConnect Performance \[page 158\]](#)

[Performance Tuning for Prepared Statements in Dynamic SQL \[page 160\]](#)

[Cursor Performance \[page 167\]](#)

## 5.1 Improve SAP jConnect Performance

Review the options to optimize the performance of an application using SAP jConnect.

- Use `TextPointer.sendData` methods to send text and image data to an SAP Adaptive Server Enterprise (SAP ASE) database.
- Create precompiled `PreparedStatement` objects for dynamic SQL statements that are used repeatedly during a session.
- Use batch updates to 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.
- 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.
- If you use `TYPE_SCROLL_INSENSITIVE` result sets, use them only when the result set is reasonably small.

## Related Information

[BigDecimal Rescaling \[page 159\]](#)

[REPEAT\\_READ Connection Property \[page 159\]](#)

[SunIoConverter Character-Set Conversion \[page 159\]](#)

[Support for Batch Updates \[page 71\]](#)

[Image Datatype \[page 74\]](#)

[Performance Tuning for Prepared Statements in Dynamic SQL \[page 160\]](#)

[TYPE\\_SCROLL\\_INSENSITIVE Result Sets in SAP jConnect \[page 68\]](#)

## 5.1.1 BigDecimal Rescaling

The JDBC 1.0 specification requires a scale factor with `getBigDecimal` method.

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 SAP 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");
    ...
}
```

## 5.1.2 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.

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.

## 5.1.3 SunIoConverter Character-Set Conversion

If you are using multibyte character sets and want to improve driver performance, use the `SunIoConverter` class provided with the SAP jConnect samples.

This converter is based on the `sun.io` classes provided by Oracle Corporation.

The `SunIoConverter` class is not a pure Java implementation of the character-set converter feature and, therefore, is not integrated with the standard SAP jConnect product, it has provided this converter class for

your reference, and you can use it with the SAP jConnect driver to improve character-set conversion performance.

#### **i Note**

Based on SAP testing, the `SunIoConverter` class improved performance on all VMs on which it was tested. However, Oracle Corporation reserves the right to remove or change the `sun.io` classes with future releases of the JDK. Therefore, this `SunIoConverter` class may not be compatible with later JDK releases.

To use the `SunIoConverter` class, you must install the SAP jConnect sample applications. Once the samples are installed, set the `CHARSET_CONVERTER_CLASS` connection property to reference the `SunIoConverter` class in the `sample2` subdirectory under your jConnect installation directory.

See the *SAP jConnect for JDBC Installation Guide* for complete instructions on installing SAP jConnect and its components, including the sample applications.

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`.

## **Related Information**

[SAP jConnect Character Set Converters \[page 47\]](#)

## **5.2 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 precompiles a dynamic statement and saves 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.

You can use SAP jConnect to optimize the performance of dynamic SQL statements on a SAP database by creating:

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

The optimal way to set the `DYNAMIC_PREPARE` connection property and create `PreparedStatement` objects can depend on whether your application needs to be portable across JDBC drivers or whether you are writing an application that allows SAP jConnect-specific extensions to JDBC.

SAP jConnect provides performance tuning features for dynamic SQL statements.

## Related Information

[Choose Prepared Statements and Stored Procedures \[page 161\]](#)  
[Prepared Statements in Portable Applications \[page 162\]](#)  
[Prepared Statements with SAP jConnect Extensions \[page 162\]](#)  
[Connection.prepareStatement \[page 164\]](#)  
[DYNAMIC\\_PREPARE Connection Property \[page 164\]](#)  
[SybConnection.prepareStatement Method \[page 165\]](#)  
[ESCAPE\\_PROCESSING\\_DEFAULT Connection Property \[page 165\]](#)  
[Optimized Batch in SAP jConnect \[page 166\]](#)  
[Choose Prepared Statements and Stored Procedures \[page 161\]](#)

### 5.2.1 Choose 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 must 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 must 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.

## 5.2.2 Prepared Statements in Portable Applications

If your application runs on databases from different vendors and you want some `PreparedStatement` objects to contain precompiled statements and others to contain uncompiled statements, use the `PreparedStatement` in portable applications.

- When you access an SAP ASE 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 because `CallableStatement` is a subclass of `PreparedStatement`, you can upcast a `CallableStatement` object to a `PreparedStatement` object, as:

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

The `PreparedStatement` object `<ps_uncomp>` is guaranteed to contain an uncompiled statement, because only `Connection.prepareStatement` is implemented to return `PreparedStatement` objects containing precompiled statements.

## 5.2.3 Prepared Statements with SAP 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 depends 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.

### Related Information

[If Most Dynamic Statements Are Executed Infrequently \[page 163\]](#)

[If Most Dynamic Statements Executed Are Multiple Times in a Session \[page 163\]](#)

### 5.2.3.1 If Most Dynamic Statements Are Executed Infrequently

Dynamic SQL statements are executed only once or twice in a session for an application.

- 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. For example:

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

### 5.2.3.2 If Most Dynamic Statements Executed Are Multiple Times in a Session

Use the `DYNAMIC_PREPARE` and `PreparedStatement` objects to execute the dynamic statements multiple times in an application in the course of a session.

- 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` or `SybConnection.prepareStatement`, with `<dynamic>` set to false. For example:

```
PreparedStatement ps_uncomp =  
    (SybConnection) conn.prepareStatement(<sql_string>, false);  
PreparedStatement ps_uncomp = Connection.prepareCall(<sql_string>);
```

## Related Information

[Prepared Statements in Portable Applications \[page 162\]](#)

## 5.2.4 Connection.PrepareStatement

SAP 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 SAP -specific applications, SAP jConnect 6.05 and later provides a `prepareStatement` method under the SAP 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.

## 5.2.5 DYNAMIC\_PREPARE Connection Property

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

- If `DYNAMIC_PREPARE` is true (the default), 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 placeholders for dynamically assigned values, and the statement needs only to be executed.
- If `DYNAMIC_PREPARE` is 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.

In this example, `DYNAMIC_PREPARE` is false to disable precompilation of dynamic SQL statements, and `props` is a `Properties` object for specifying connection properties.

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

When `DYNAMIC_PREPARE` is true:

- 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 cannot precompile and save a statement sent to it through `Connection.prepareStatement`, SAP jConnect 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, explicitly close every `PreparedStatement` object after its last use to prevent prepared statements from accumulating in server memory during a session and slowing performance.

## 5.2.6 SybConnection.PrepareStatement Method

Use the `SybConnection.prepareStatement` extension method to return dynamic SQL statements in `PreparedStatement` objects.

If your application allows SAP jConnect-specific extensions to JDBC:

```
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.

This 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 this example, the connection object `<conn>` is cast to a `SybConnection` object to allow the use of `SybConnection.prepareStatement`. The SQL string passed to `SybConnection.prepareStatement` is precompiled in the database, even if the connection property `DYNAMIC_PREPARE` is false.

If the database generates an error because it cannot to precompile a statement sent to it through `SybConnection.prepareStatement`, SAP 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.

## 5.2.7 ESCAPE\_PROCESSING\_DEFAULT Connection Property

By default, SAP 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, set this connection property to false to circumvent this parsing. This may give a slight performance benefit.



## 5.2.8 Optimized Batch in SAP jConnect

SAP jConnect implements an internal algorithm to speed up batch operations for `PreparedStatement` objects.

This algorithm is invoked when the `HOMOGENEOUS_BATCH` connection property is true.

### i Note

Homogeneous batching is available only when your client application is connected to a server that supports this feature. SAP Adaptive Server Enterprise (SAP ASE) 15.7 introduces support for homogeneous batching.

This example illustrates a `PreparedStatement` batching operation using the `addBatch` and `executeBatch` methods:

```
String sql = "update members set lastname = ? where member_id = ?";
prep_stmt = connection.prepareStatement(sql);
prep_stmt.setString(1, "Forrester");
prep_stmt.setLong(2, 45129);
prep_stmt.addBatch();
prep_stmt.setString(1, "Robinson");
prep_stmt.setLong(2, 45130);
prep_stmt.addBatch();
prep_stmt.setString(1, "Servo");
prep_stmt.setLong(2, 45131);
prep_stmt.addBatch();
prep_stmt.executeBatch();
```

where `connection` is a connection instance, `prep_stmt` is a prepared statement instance, and `?` denotes parameter placeholders for the prepared statement.

## Related Information

[Homogeneous Batch with Large Object \(LOB\) Columns \[page 166\]](#)

### 5.2.8.1 Homogeneous Batch with Large Object (LOB) Columns

When the `HOMOGENEOUS_BATCH` and `ENABLE_LOB_LOCATORS` properties are true, the client application cannot mix LOB and non-LOB prepared statement setter methods in the same batch.

For example, this is invalid:

```
String sql = "update members SET catchphrase = ? WHERE member_id = ?";
prep_stmt = connection.prepareStatement(sql);
prep_stmt.setString(1, "Push the button, Kal-El!");
prep_stmt.setLong(2, 45129);
prep_stmt.addBatch();
Clob myclob = con.createClob();
```

```
myclob.setString(1, "Hi-Jor-El!");
prep_stmt.setClob(1, myclob);
prep_stmt.setLong(2, 45130);
prep_stmt.addBatch();
pstmt.executeBatch();
```

where catchphrase is a column of type text. This code fails because the `setString` method and the `setClob` method are used in the same batch for the same column.

## 5.3 Cursor Performance

When you use the `Statement.setCursorName` method or the `setFetchSize( )` method in the `SybCursorResultSet` class, SAP jConnect creates a cursor in the database.

Other methods cause SAP jConnect to open, fetch, and update a cursor.

SAP jConnect creates and manipulates cursors either by sending SQL statements to the database or by encoding cursor commands as tokens within the 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, SAP SQL Anywhere databases do not support language cursors.

In SAP jConnect, the default condition is for all cursors to be protocol cursors. However, the `LANGUAGE_CURSOR` connection property lets you use language commands in the database to create and manipulate cursors.

### Related Information

[LANGUAGE\\_CURSOR Connection Property \[page 167\]](#)

### 5.3.1 LANGUAGE\_CURSOR Connection Property

`LANGUAGE_CURSOR` is a Boolean-valued connection property in SAP jConnect that allows you to determine whether cursors are created as protocol cursors or language cursors.

- If `LANGUAGE_CURSOR` is false (the default), all cursors created during a session are protocol cursors, which provide better performance. SAP jConnect creates and manipulates the cursors by sending cursor commands as tokens in the TDS protocol.
- If `LANGUAGE_CURSOR` is true, all cursors created during a session are language cursors. SAP 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.

## 6 Migrating SAP jConnect Applications

Review instructions to migrate applications from SAP jConnect 5.x ,6.x, 7.x, and SAP jConnect 16.x.

### Related Information

[Migrating Applications to SAP jConnect 16.x \[page 168\]](#)

[Change SAP jConnect Extensions \[page 169\]](#)

### 6.1 Migrating Applications to SAP jConnect 16.x

Review the instructions to migrate applications to SAP jConnect 16.x.

#### Procedure

1. If your code uses SAP jConnect extensions, or if you explicitly import any SAP jConnect class in your code, change package import statements as needed.

For example, change import statements such as:

```
import com.sybase.jdbc.*
```

and:

```
import com.sybase.jdbc2.jdbc.*
```

to:

```
import com.sybase.jdbcx.*
```

2. Set JDBC\_HOME to the top directory of the SAP jConnect driver you installed:

```
JDBC_HOME=jConnect-16_0
```

3. Change your CLASSPATH environment variable to reflect the new installation; it must include:

```
JDBC_HOME/classes/jconn4.jar
```

4. Change the source code where the driver is loaded, and recompile the application to use the new driver:

```
Class.forName("com.sybase.jdbc4.jdbc.SybDriver");
```

5. Verify that the SAP jConnect 16.0 driver is the first SAP jConnect driver specified in your CLASSPATH environment variable.

## Related Information

[Change SAP jConnect Extensions \[page 169\]](#)

## 6.2 Change SAP jConnect Extensions

SAP jConnect version 4.1 and later include the package `com.sybase.jdbcx` that contains all of the SAP jConnect extensions to JDBC.

In versions of SAP jConnect earlier than 4.1, these extensions were available in the `com.sybase.jdbc` and `com.sybase.utils` packages.

The `com.sybase.jdbcx` package provides a consistent interface across different versions of SAP jConnect. All of the SAP jConnect extensions are defined as Java interfaces, which allow the underlying implementations to change without affecting applications built using these interfaces.

When you develop new applications that use SAP jConnect extensions, use `com.sybase.jdbcx`. The interfaces in this package allow you to upgrade applications to versions of SAP jConnect later than 4.0 with minimal changes.

Some of the SAP jConnect extensions have been changed to accommodate the `com.sybase.jdbcx` interface.

## Related Information

[Extension Change Example \[page 169\]](#)

[Method Names \[page 170\]](#)

[Debug Class \[page 170\]](#)

### 6.2.1 Extension Change Example

Review the code differences if an application uses the `SybMessageHandler`.

- SAP 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());
```

- SAP jConnect 6.0 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 SAP jConnect for more examples of how to use SAP jConnect extensions.

## 6.2.2 Method Names

Review the list of renamed methods in the interface.

Method Name Changes

Actual Method Name	Version 4.0 and Earlier	Version 4.0 and Later
SybConnection	getCapture ( )	createCapture ( )
SybConnection	setMessageHandler ( )	setSybMessageHandler ( )
SybConnection	getMessageHandler ( )	getSybMessageHandler ( )
SybStatement	setMessageHandler ( )	setSybMessageHandler ( )
SybStatement	getMessageHandler ( )	getSybMessageHandler ( )

## 6.2.3 Debug Class

Direct static references to the `Debug` class are no longer supported, but exist in deprecated form in the `com.sybase.utils` package.

Use SAP 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.jdbc4.jdbc.SybDriver") newInstance();
Debug sybDebug = sybDriver.getDebug();
sybDebug.debug(true, "ALL", System.out);
```

A complete list of SAP jConnect extensions is in the SAP jConnect Javadoc documentation located in the `docs/` directory of your SAP jConnect installation directory.

# 7 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 SSL protocol, SAP jConnect includes a Java servlet that you can install on any Web server that supports the `javax.servlet` interfaces. This servlet enables SAP jConnect to support encryption using the Web server as the gateway.

## i Note

SAP jConnect includes support for SSL on the client system.

## Related Information

[TDS tunnelling \[page 172\]](#)

[Configure SAP jConnect and Gateways \[page 173\]](#)

[Usage Requirements \[page 175\]](#)

[TDS-Tunnelling Servlet \[page 177\]](#)

[Implement Custom SSL Socket Plug-ins \[page 126\]](#)

## 7.1 TDS tunnelling

SAP jConnect uses TDS to communicate with database servers. Requests from a client to a back-end server that go through the gateway contain TDS in the body of the request.

HTTP-tunnelled TDS is useful for forwarding requests. 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, SAP 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 SAP jConnect to support encrypted communication between two systems; for example, a SAP jConnect client with the `CONNECT_PROTOCOL` connection property set to `https` can connect to a Web server running the TDS-tunnelling servlet.

## 7.2 Configure SAP jConnect and Gateways

There are several options for setting up your Web servers and SAP Adaptive Server Enterprise (SAP ASE) servers to install the SAP jConnect driver and to use a gateway with the TDS-tunnelling servlet.

### Related Information

[Web Server and SAP ASE on One Host \[page 173\]](#)

[Dedicated JDBC Web Server and SAP ASE on One Host \[page 173\]](#)


[Web Server and SAP ASE on Separate Hosts \[page 174\]](#)

[Connect to Server Through Firewall \[page 175\]](#)

### 7.2.1 Web Server and SAP ASE on One Host

In a two-tier configuration, the Web server and SAP Adaptive Server Enterprise (SAP ASE) are both installed on the same host.

- Install SAP jConnect on the Web server host.
- No gateway is required.



Web Server and SAP ASE on One Host

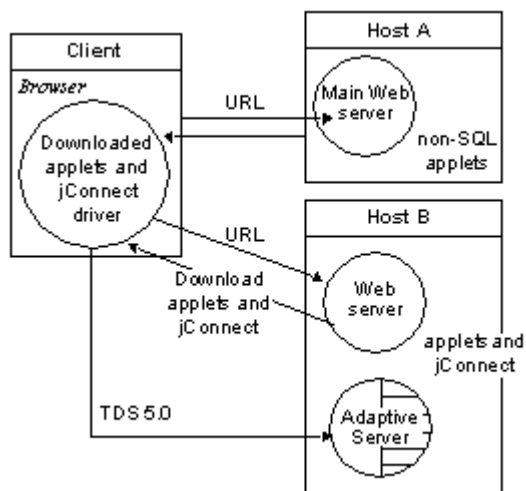
### 7.2.2 Dedicated JDBC Web Server and SAP ASE on One Host

In a single host configuration, you have a separate host for your main Web server.

A second host is shared by a Web server specifically for SAP Adaptive Server Enterprise (SAP ASE) access and the SAP ASE. Links from the main server send requests requiring SQL access to the dedicated Web server.

- Install SAP jConnect on the second (SAP ASE) host.
- No gateway is required on the second (SAP ASE) host.



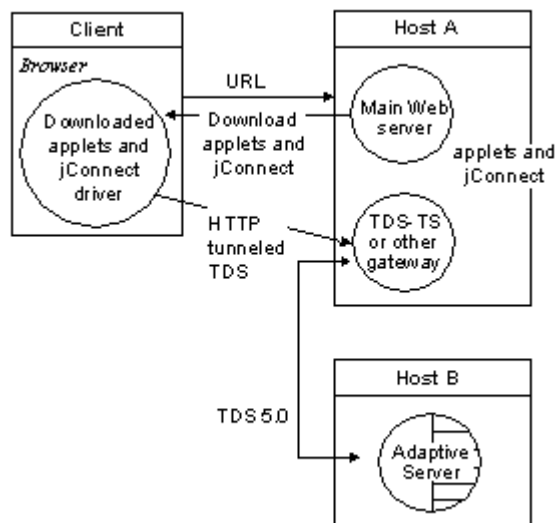


Dedicated JDBC Web Server and SAP ASE on One Host

## 7.2.3 Web Server and SAP ASE on Separate Hosts

In a three-tier configuration, the SAP Adaptive Server Enterprise (SAP ASE) is on a separate host than the Web server. SAP jConnect requires a gateway to act as a proxy to the SAP ASE.

- Install SAP jConnect on the Web server host.
- Install a TDS-tunnelling servlet or a different gateway.




Web Server and SAP ASE on Separate Hosts

## 7.2.4 Connect to Server Through Firewall

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 SAP jConnect on the Web server host.
- Requires a Web server that supports the `javax.servlet` interfaces.



Connecting to a Server Through a Firewall

## 7.3 Usage Requirements

Review the usage requirements for Web server gateways.

### Related Information

[Viewing the Index.html File \[page 176\]](#)

[Running Sample Applet \[page 176\]](#)

[Modifying Applet Screen Dimensions \[page 177\]](#)

## 7.3.1 Viewing the Index.html File

Use your Web browser to view the `index.html` file in your SAP jConnect installation directory. `index.html` provides links to the SAP jConnect documentation and sample code.

### Context

If you use Netscape on the same machine where you have installed SAP jConnect, be sure that your browser does not have access to your CLASSPATH environment variable. See *Restrictions on Setting CLASSPATH When You Use Netscape* in the *SAP jConnect for JDBC Installation Guide and Release Bulletin*.

### Procedure

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.

## 7.3.2 Running Sample Applet

Instructions to execute the sample applet in SAP jConnect.

### Procedure

1. Click **Run Sample JDBC Applets**.
2. In the Executable Samples table, locate `Isql.java` and click **Run** at the end of the row.

## Results

The sample `Isql.java` applet prompts for a simple query on a sample database and displays the results. The applet displays a default SAP Adaptive Server Enterprise (SAP ASE) host name, port number, user name (`<guest>`), password (`<sybase>`), database, and query. Using the default values, the applet connects to the SAP jConnect demonstration database, and returns results after you click **Go**.

### 7.3.3 Modifying Applet Screen Dimensions

On UNIX platforms, if the applet does not appear as expected, you can modify the applet screen dimensions.

#### Procedure

1. Use a text editor to open `$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.

## 7.4 TDS-Tunnelling Servlet

To use the TDS-tunnelling servlet, you need a Web server that supports the `javax.servlet` interfaces, such as the Oracle Corporation Java Web server.

When you install the Web server, include the SAP 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 of 1234.

<b>Client to gateway login request</b>	<p>No session ID required.</p> <ul style="list-style-type: none"> <li>• Query– POST/tds?ServerHost=dbserver&amp;ServerPort=1234&amp;Operation=more HTTP/1.0</li> <li>• Header – Content-Length: 605</li> <li>• Content (TDS) – Login request</li> </ul>
<b>Gateway to client</b>	<p>Header contains session ID assigned by the TDS servlet.</p> <ul style="list-style-type: none"> <li>• Query– 200 SUCCESS HTTP/1.0</li> <li>• Header – Content-Length: 210 TDS-Session: TDS00245817298274292</li> <li>• Content (TDS) – Login acknowledgment EED</li> </ul>
<b>Client to gateway</b>	<p>Headers for all subsequent requests contain the session ID.</p> <ul style="list-style-type: none"> <li>• Query– POST/tds?TDS-Session=TDS00245817298274292&amp;Operation=more HTTP/1.0</li> <li>• Header – Content-Length: 32</li> <li>• Content (TDS) – Query “SELECT * from authors”</li> </ul>
<b>Gateway to client</b>	<p>Headers for all subsequent responses contain the session ID</p> <ul style="list-style-type: none"> <li>• Query– 200 SUCCESS HTTP/1.0</li> <li>• Header – Content-Length: 2048 TDS-Session: TDS00245817298274292</li> <li>• Content (TDS) – Row format and some rows from query response</li> </ul>

## Related Information

[Reviewing Requirements \[page 178\]](#)

[Installing and Setting Servlet Arguments \[page 179\]](#)

[Invoking the Servlet \[page 180\]](#)

[Tracking Active TDS Sessions \[page 180\]](#)

[Terminating TDS Sessions \[page 180\]](#)

[Resuming a TDS Session \[page 181\]](#)

## 7.4.1 Reviewing Requirements

To use SAP jConnect servlet for TDS-tunnelling, you must have a Web server that supports the `javax.servlet` interface.

## Context

To install the server, follow the instructions that are provided with the Java servlet.

## 7.4.2 Installing and Setting Servlet Arguments

SAP jConnect installation includes a `gateway2` subdirectory under the `classes` directory. The subdirectory contains files required for the TDS-tunnelling servlet.

### Procedure

1. Copy the SAP jConnect `gateway` package to a `gateway2` subdirectory under the `servlets` directory of your Web server.

After you have copied the servlets, activate the servlets by following the instructions for your Web server.

2. Add the servlet to the Web server and, to customize performance set the optional arguments:
  - `SkipDoneProc` [true|false] – databases often return row count information while 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 removes this extra information from responses, which reduces network usage and processing requirements on the client. This is particularly useful when using HTTPS/SSL, because the unwanted data is not encrypted/decrypted.
  - `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 SAP 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.

### Related Information

[Debugging with SAP jConnect \[page 146\]](#)

## 7.4.3 Invoking the Servlet

SAP jConnect determines when to use the gateway where the TDS-tunnelling servlet is installed based on the path extension of the `proxy` connection property.

### Context

SAP 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>
```

SAP 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 the servlet alias list for your Web server.

## 7.4.4 Tracking Active TDS Sessions

View information about active TDS sessions, including the server connections, for each session.

### ❖ Example

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.

## 7.4.5 Terminating TDS Sessions

To terminate a TDS session, use the URL defined in any active TDS session.

### Context

Select an active session from the list of sessions on the first page, then click **Terminate This Session**.

## 7.4.6 Resuming a TDS Session

When you specify a `SESSION_ID`, SAP jConnect skips the login phase of the protocol and resumes the connection with the gateway using the designated session ID.

### Context

Set the `SESSION_ID` connection property so that, if necessary, you can resume an existing open connection.

If the session ID you specified does not exist on the servlet, SAP jConnect throws a SQL exception the first time you attempt to use the connection.



# 8 SAP jConnect Sample Programs

Review SAP jConnect sample programs.

## Related Information

[Running IsqlApp \[page 182\]](#)

## 8.1 Running IsqlApp

IsqlApp allows you to issue `isql` commands from the command line, and run SAP jConnect sample programs.

## Context

The syntax for IsqlApp:

```
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>]
        [-K <service_principal_name>]
        [-F <JAAS_login_config_file_path>]
        [-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.

Parameter	Description
-G	<p>The gateway address. For HTTP, the URL is: <code>http://&lt;host:port&gt;</code>.</p> <p>To use HTTPS, the URL is <code>https://&lt;host&gt;:&lt;port&gt;/&lt;servlet_alias&gt;</code>.</p>
-p	Whether to use HTTP or HTTPS.
-D	<p>Turns on debugging for all classes or for just the ones you specify, separated by a comma. For example:</p> <p>-D ALL – displays debugging output for all classes.</p> <p>-D SybConnection, Tds – displays debugging output only for the <code>SybConnection</code> and <code>Tds</code> classes.</p>
-v	Turns on verbose output for display or printing.
-I	<p>Causes <code>IsqlApp</code> to take commands from a file instead of the keyboard.</p> <p>After the parameter, specify the name of the file to use for the <code>IsqlApp</code> input. The file must contain command terminators ("go" by default).</p>
-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	<p>Specifies the character set for strings passed through TDS.</p> <p>If you do not specify a character set, <code>IsqlApp</code> uses the default charset of the server.</p>
-L	Specifies the language in which to display error messages returned from the server, and for SAP jConnect messages.
-K	<p>Indicates the user wants to make a Kerberos login to SAP Adaptive Server. This parameter sets the service principal name. For example:</p> <pre>-K myASE</pre> <p>where the service principal name for your server is <code>myASE</code>.</p>
-F	<p>Specifies the path to the JAAS login configuration file. You must set this property if you use the <code>-K</code> option. For example:</p> <pre>-F /foo/bar/exampleLogin.conf</pre> <p>See the <code>ConnectKerberos.java</code> sample in the <code>sample2</code> directory of your SAP jConnect installation.</p>

Parameter	Description
-T	When this parameter is set, SAP jConnect assumes that an application is trying to resume communication on an existing TDS session held open by the TDS-tunnelling gateway. SAP jConnect skips the login negotiations and forwards all requests from the application to the specified session ID.
-V	Enables the use of version-specific characteristics.

You must enter a space after each option flag.

To obtain a full description of the command line options, enter:

```
java IsqlApp -help
```

This 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
```

An applet that provides GUI access to `isql` commands is available as: `$JDBC_HOME/sample2/gateway.html` (UNIX) `%JDBC_HOME%\sample2\gateway.html` (Windows).

## Related Information

[Security \[page 125\]](#)

[JCONNECT\\_VERSION Connection Property \[page 9\]](#)

## 9 SAP jConnect Sample Programs and Code

SAP jConnect includes several sample programs that are intended to help you understand how SAP jConnect works with various JDBC classes and methods.

### Related Information

[Sample Applications \[page 185\]](#)

[Sample Code \[page 186\]](#)

### 9.1 Sample Applications

When you install SAP jConnect, you can also install sample programs, which include the source code so that you can review how SAP jConnect implements various JDBC classes and methods.

See the *SAP jConnect for JDBC Installation Guide* for complete instructions for installing the sample programs.

#### **i** Note

The SAP jConnect sample programs are intended for demonstration purposes only.

The sample programs are installed in the `sample2` subdirectory under your SAP 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.

### Related Information

[Running the Sample Applets \[page 186\]](#)

[Running the Sample Programs with SAP SQL Anywhere \[page 186\]](#)

## 9.1.1 Running the Sample Applets

You can run some of the sample programs as applets in your Web browser, enabling you to view the source code while you review the output results.

### Context

To run the sample programs as applets, enter `http://localhost:8000/sample2/index.html` on a Web browser to start the Web server gateway.

## 9.1.2 Running the Sample Programs with SAP SQL Anywhere

All of the sample programs are compatible with SAP Adaptive Server, but only a limited number are compatible with SAP SQL Anywhere.

### Context

Refer to `index.html` in the `sample2` subdirectory for a current list of the sample programs that are compatible with SAP SQL Anywhere.

To run the sample programs that are available for SAP SQL Anywhere, you must install the `pubs2_any.sql` script on your SAP SQL 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/pubs2_any.sql -c go
```

## 9.2 Sample Code

Review the sample code that illustrates how to invoke the SAP 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.jdbc4.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);  
}  
}
```

# 10 SQL Exception and Warning Messages

Review the SQL exception and warning messages that you may encounter when using SAP 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>Action: Use the devclasses debug classes to determine the reason for this message, and report the problem to SAP Technical Support.</p>
010CP	<p>AutoCommit option has changed to true. All pending statements on this transaction (if any) are committed.</p>
010DF	<p>Attempt to set database at login failed. Error message:_____.</p> <p>Description: SAP 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 you are connecting to SAP SQL Anywhere, use the <code>SERVICENAME</code> 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." SAP 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: Be sure your application defines the connection property only once. However, you may want your 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>Action: Reconfigure the server to support high availability failover or do not set <code>REQUEST_HA_SESSION</code> to true</p>
010HD	<p>SAP Adaptive Server high-availability failover is not supported by this type of database server.</p> <p>Action: Connect only to database servers that support high availability failover.</p>
010HN	<p>The client did not specify a <code>SERVICE_PRINCIPAL_NAME</code> Connection property. Therefore, jConnect is using the hostname of _____ as the service principal name</p> <p>Action: Explicitly specify a service principal name using the connection property.</p>



SQL state	Message/Description/Action
010HT	<p>Hostname property truncated, maximum length is 30.</p> <p>Action: No action is necessary, since this is just a warning to indicate that SAP 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>
010KF	<p>The server rejected your Kerberos login attempt. Most likely, this was because of a Generic Security Services (GSS)exception. Please check your Kerberos environment and configuration.</p> <p>Action: Check your Kerberos environment, and make sure that you authenticated properly to the KDC. See <a href="#">Security [page 125]</a> for more information.</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: Be sure that stored procedures for providing metadata are installed on the server. See <i>Installing Stored Procedures</i> in the <i>SAP 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, 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 SAP jConnect from returning this warning, and possibly see performance improvement, by setting the DISABLE_UNPROCESSED_PARAM_WARNINGS connection property to true.</p>
010P6	<p>A row was received and ignored.</p> <p>Description: An unexpected object of type 0xD1 was encountered in the result set being processed and has been ignored.</p> <p>Action: Check the query that generated the result set and correct as required.</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 file 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 file names specified does not exist on the server, the above error message results.</p> <p>Action: Verify that every .jar file mentioned in the PRELOAD_JARS connection property for your application exists and is accessible on the server.</p>

SQL state	Message/Description/Action
010PO	<p>Property <code>LITERAL_PARAM</code> has been reset to false because <code>DYNAMIC_PREPARE</code> was set to true.</p> <p>Description: To use precompiled dynamic statements, allow for parameters to be sent to those statements (if the statements take parameters). Setting <code>LITERAL_PARAMS</code> 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 <code>LITERAL_PARAMS</code> to true when you want to use dynamic SQL. See <a href="#">Performance Tuning for Prepared Statements in Dynamic SQL [page 160]</a>.</p>
010RC	<p>The requested <code>ResultSet</code> type and concurrency is not supported. They have been converted.</p> <p>Description: See <a href="#">Use Cursors with Result Sets [page 60]</a> for more information about what result set types and concurrencies are available in SAP jConnect</p> <p>Action: Request a supported type and concurrency combination for the result set.</p>
010SJ	<p>Metadata accessor information was not found on this database. Please install the required tables as mentioned in the jConnect 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 jConnect see <i>Installing Stored Procedures</i> in the <i>jConnect for JDBC Installation Guide</i>. If you do not need metadata, set the <code>USE_METADATA</code> 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: 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 SAP jConnect. See <i>Installing Stored Procedures</i> in the <i>SAP jConnect 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: SAP Adaptive Server Enterprise (SAP ASE) versions 11.9.2 and lower had an issue that caused them to reject logins from clients that requested capabilities that the servers did not have. This warning indicates that SAP jConnect has detected this condition and is retrying the connection using the greatest number of capabilities that the server can accept. When SAP jConnect encounters this bug, it attempts to connect to the server twice.</p> <p>Action: Clients can safely ignore this warning, but to eliminate the warning and ensure that SAP jConnect makes only one connection attempt, they can set the <code>ELIMINATE_010SM</code> connection property to true. Do not set this property to true when connecting to SAP 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 <code>PROTOCOL_CAPTURE</code> 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 <code>CONNECTION_FAILOVER</code> connection property is set to true, and SAP jConnect cannot to connect to one of the database servers in the list of servers to which to connect. Therefore, SAP jConnect now tries to connect to the next server in the list.</p> <p>Action: No action is required, as long as SAP jConnect can to connect to another database server. However, determine why SAP jConnect could not 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 SAP jConnect, and has responded with a different character set. SAP jConnect accepts the change, and performs 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 <code>CHARSET</code> 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: When this occurs, the only characters that are guaranteed to translate properly are the first 127 ASCII characters. Therefore, SAP jConnect reverts to 7-bit ASCII. The message is strictly informational and has no further consequences.</p> <p>Action: Install the stored procedures for returning metadata that comes with SAP jConnect. See <i>Installing Stored Procedures</i> in the <i>SAP jConnect for JDBC Installation Guide</i>.</p>
010UF	<p>Attempt to execute use database command failed. Error message:_____.</p> <p>Description: SAP jConnect cannot connect to the database specified in the connection URL. Two possibilities are:</p> <ul style="list-style-type: none"> <li>• The name was entered incorrectly in the URL.</li> <li>• <code>USE_METADATA</code> is true (the default ), but the stored procedures for returning metadata are not installed. As a result, SAP jConnect tried to execute the <code>use &lt;database &gt;</code> command with the database in the URL, but the command failed. This may be because you attempted to access a SQL Anywhere databases, which does not support the <code>use &lt;database &gt;</code> command.</li> </ul> <p>Action: Make sure the database name in the URL is correct. Make sure the stored procedures for returning metadata are installed on the server. See <i>Installing Stored Procedures</i> in the <i>SAP jConnect for JDBC Installation Guide</i> and <i>SAP jConnect for JDBC 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 SAP jConnect driver connection properties.</p>
0100V	<p>The version of TDS protocol being used is too old. Version: _____.</p> <p>Description: SAP jConnect requires version 5.0 or later.</p> <p>Action: Use a server that supports the required version of TDS. See the <i>System Requirements</i> in the <i>SAP jConnect Installation Guide</i>.</p>
01S07	<p>Adaptive Server may round or truncate nanosecond values.</p> <p>Description: A time value more precise than 1/300th of a second encountered. Because SAP Adaptive Server does not support such precision, jConnect rejected the value.</p> <p>Action: Make sure that time values are of precision up to 1/300th of a second.</p>

SQL state	Message/Description/Action
01S08	<p>This connection has been enlisted in a Global transaction. All pending statements on the current local transaction (if any) have been rolled back.</p> <p>Description: SAP jConnect issues rollback to clear out any current local transactions. This occurs when a global transaction has been enlisted, after issuing <code>XAResource.start()</code>.</p> <p>Action: Either commit or roll back active local transactions before issuing the <code>XAResource.start()</code> method.</p>
01S09	<p>The local transaction method _____ cannot be used while a global transaction is active on this connection.</p> <p>Description: An example of a local operation is calling the <code>commit()</code> method on the connection. Other operations that cannot be used include: <code>rollback()</code>, <code>rollback(Savepoint)</code>, <code>setSavepoint()</code>, <code>setSavepoint(String)</code>, <code>releaseSavepoint(Savepoint)</code>, and <code>setAutoCommit()</code>.</p> <p>Action: Keep local transactions separate from global transactions. Complete all local transactions and their operations before starting the global transaction.</p>
01S10	<p>The local transaction method _____ cannot be used on a pre-System 12 XAConnection.</p> <p>Description: You have used a local transaction method that does not work with versions earlier than SAP SQL Anywhere version 12.</p> <p>Action: Do not use the method.</p>
01S11	<p>WARNING: Your data might be truncated.</p> <p>Description: The user-specified length of a stream or LOB is greater than the limit (Integer.MAX_VALUE) in a <code>ResultSet.updateXXX</code> method.</p> <p>Action: Make sure that the length is within the limit.</p>
01S12	<p>Unable to continue with HOMOGENEOUS_BATCH protocol, falling back to normal batching.</p> <p>Description: When <code>DYNAMIC_PREPARE</code> is set to false, SAP Adaptive Server does not send parameter metadata. When <code>HOMOGENEOUS_BATCH</code> is true, SAP jConnect requires this information for optimization. Thus, jConnect reverts to normal batching.</p> <p>Action: Use optimized batching (<code>HOMOGENEOUS_BATCH</code> set to true) with precompiled dynamic SQL prepared statements only (<code>DYNAMIC_PREPARE</code> set to true).</p>
01S13	<p>Connected Adaptive Server server does not support the set option 'logbulkcopy' needed for logging BCP. Falling back to normal bulk load without logging which is equivalent to setting <code>ENABLE_BULK_LOAD=BCP</code>.</p> <p>Description: If the SAP Adaptive Server is earlier than 15.7 ESD #1 and does not support logged bulk loading, SAP jConnect reverts to normal batching.</p> <p>Action: Use <code>ENABLE_BULK_LOAD=LOG_BCP</code> setting only with SAP Adaptive Server 15.7 ESD #1 or later.</p>

SQL state	Message/Description/Action
01ZZZ	<p>Error code 4022:</p> <p>Password has expired Please set the NEWPASSWORD property with the new password or use sp_password to change passwords.</p> <p>Action: Reset the password for connecting to SAP Adaptive Server.</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>
JZ003	<p>Incorrect URL format. URL: _____.            Action: Verify the URL format. See <a href="#">URL connection property parameters [page 38]</a>.</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> <ul style="list-style-type: none"> <li>• The PROXY format for the Cascade proxy is: &lt;ip_address&gt;:&lt;port_number&gt;</li> <li>• The PROXY format for the TDS tunnelling servlet is: http[s]://&lt;host&gt;:&lt;port&gt;/&lt;tunneling_servlet_alias&gt;</li> </ul>
JZ004	<p>User name property missing in DriverManager.getConnection(..., Properties)</p> <p>Action: Provide the required user property.</p>
JZ006	<p>Caught IOException: _____.            Description: An unexpected I/O error is 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, SAP jConnect chains 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.            Action: Examine the text of the original I/O exception message, and proceed from there.</p>
JZ008	<p>Invalid column index value _____.            Description: You have requested a column index value of less than 1 or greater than the highest available value.            Action: Verify the call to the getXXX method and the text of the original query, or call rs.next.</p>

SQL state	Message/Description/Action
JZ009	<p>Error encountered in conversion. Error message: ____.</p> <p>Description: Possibilities include:</p> <ul style="list-style-type: none"> <li>• An attempt to convert between two incompatible datatypes, such as date to int.</li> <li>• An attempt to convert a string containing a nonnumeric character to a numeric type.</li> <li>• A formatting error, such as an incorrectly formatted time/date string.</li> </ul> <p>Action: Make sure that the JDBC specification supports the attempted type conversion. Make sure that strings are correctly formatted. If a string contains nonnumeric characters, do not attempt to convert it to a numeric type.</p>
JZ00A	<p>Invalid precision and scale specified for a numeric value.</p> <p>Description: When using the <code>setBigDecimal</code> method, the <code>BigDecimal</code> value is set to either a precision value of less than 1, a negative scale value, a precision less than the scale value, or precision value greater than 127.</p> <p>Action: Examine the query and correct to specify a legal precision/scale value.</p>
JZ00B	<p>Numeric overflow.</p> <p>Description: You tried to send a <code>BigInteger</code> as a TDS numeric, and the value was too large, or you tried to send a Java <code>long</code> as an <code>int</code> and the value was too large.</p> <p>Action: These values cannot be stored in SAP jConnect. For <code>long</code>, consider using a SAP jConnect numeric. There is no solution for <code>Bignum</code>.</p>
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: Use a separate statement to delete or update a cursor. See <a href="#">Use Cursors with Result Sets [page 60]</a>.</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: Call <code>updateXX</code> methods before calling <code>updateRow</code>.</p>

SQL state	Message/Description/Action
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 <math>\geq 0</math>.</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 SQLWarnings 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 CLASSPATH variable.</p>
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 SAP technical support. Wrong access type for connection property ____.</p> <p>Action: Contact SAP 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>
JZ015	<p>Illegal value set for the <code>GSSMANAGER_CLASS</code> connection property. The property value must be a String or an Object that extends <code>org.ietf.jgss.GSSManager</code>.</p> <p>Action: Check the value to which you set the <code>GSSMANAGER_CLASS</code> property.</p>



SQL state	Message/Description/Action
JZ017	<p>Savepoint is not valid.</p> <p>Description: You have specified a nonexistent savepoint for rollback or release.</p> <p>Action: Examine the query and correct to specify an existing savepoint.</p>
JZ018	<p>This method can not be applied to this type of savepoint.</p> <p>Description: The <code>getSavepointId()</code> method does not work on named savepoints (which have no ID), and the <code>getSavepointName()</code> method does not work on unnamed savepoints (which have no name).</p> <p>Action: Examine the query and correct.</p>
JZ019	<p>Error obtaining SERVERNAME : ____.</p> <p>Description: The URL set using <code>jdbc: :jndi:file</code> does not specify either the <code>sql.ini</code> file (Windows) or the <code>interfaces</code> file (UNIX) or a server name.</p> <p>Action: Examine the URL command and correct.</p>
JZ021	<p>The Specified ____ file not found.</p> <p>Description: The <code>sql.ini</code> file (Windows) or the <code>interfaces</code> file (UNIX) specified in the connection URL is not found.</p> <p>Action: Check the connection URL and correct.</p>
JZ022	<p>The Specified ____ file has an unknown format.</p> <p>Description: The connection URL string in the <code>sql.ini</code> file (Windows) or the <code>interfaces</code> file (UNIX) is not in the correct format.</p> <p>Action: Check the connection URL string and correct.</p>
JZ024	<p>The Specified Server : ____ has no entry in the <code>interfaces/sql.ini</code> file : ____.</p> <p>Action: Check the connection URL string and correct.</p>
JZ025	<p>The TLI format for Specified Server in <code>interfaces/sql.ini</code> is Invalid.</p> <p>Action: Check the settings and correct.</p>
JZ026	<p>The Specified Protocol : ____ for Server : ____ in <code>interfaces/sql.ini</code> file : ____ is not Supported.</p> <p>Description: A protocol other than TLI, TCP, and NLWNSCK is specified in the <code>sql.ini</code> file (Windows) or the <code>interfaces</code> file (UNIX).</p> <p>Action: Specify a supported protocol.</p>
JZ027	<p>The Specified SECMECH entrys: ____ for Server : ____ in <code>interfaces/sql.ini</code> file : ____ are not Supported.</p> <p>Description: An invalid value is specified in the Kerberos connection URL.</p> <p>Action: Examine the URL and correct.</p>

SQL state	Message/Description/Action
JZ028	<p>Illegal value set for JCE_PROVIDER_CLASS connection property. The property value must be a fully qualified provider class name passed as a String or an instance of java.security.Provider.</p> <p>Action: Specify a legal value.</p>
JZ029	<p>Error looking up address for ALTERNATE_SERVER_NAME _____. (____).</p> <p>Description: SAP jConnect cannot locate the server specified with the ALTERNATE_SERVER_NAME property using the SAP SQL Anywhere UDP discovery protocol.</p> <p>Action: Check the server name specified with the ALTERNATE_SERVER_NAME connection property and correct.</p>
JZ030	<p>The method _____ is not supported.</p>
JZ031	<p>Failed to unwrap the object of _____.</p> <p>Description: SAP jConnect cannot unwrap the object of a custom class because the custom class is not in the classpath.</p> <p>Action: Add the class to classpath.</p>
JZ032	<p>A Date or Timestamp parameter exceeds the BigDateTime/BigTime range. The server can only support BigDateTime values between 0001/01/01 12:00:00:000000AM to 9999/12/31 11:59:59.999999PM or BigTime values between 12:00:00:000000AM to 11:59:59.999999PM.</p>
JZ033	<p>Unknown Blob type returned by the server.</p> <p>Description: SAP jConnect cannot map the SAP Adaptive Server datatype of the column to a BLOB datatype.</p> <p>Action: Make sure that the SAP Adaptive Server column is convertible to a BLOB datatype.</p>
JZ034	<p>The connected server is not capable of handling Large Objects [LOB].</p> <p>Action: Use the regular stream methods to access LOB.</p>
JZ035	<p>To handle Large object [LOB], please set connection property 'ENABLE_LOB_LOCATOR' to true.</p>
JZ036	<p>Reference to this Large Object [LOB] is no longer valid in database. Check if you have called free() or check if transaction ended.</p>
JZ037	<p>Value of offset/position/start should be in the range [1, len] where len is the length of Large Object[LOB].</p>
JZ038	<p>Length of the object should be &gt;= 0.</p>
JZ040	<p>_____ operation failed. The _____ has been closed already.</p> <p>Description: The read (write) operation has failed as input stream or LOB reader (output stream or LOB writer) has already closed</p> <p>Action: Check the application to locate the reason of the conflict and correct.</p>

SQL state	Message/Description/Action
JZ041	<p>_____ operation failed on the _____.</p> <p>Description: The read(write)(available()) operation has failed as input stream or reader (output stream or writer)(input stream) has already closed.</p> <p>Action: Check the application to locate the reason of the conflict and correct.</p>
JZ042	<p>Large Object setters can not be mixed with other setters when ENABLE_LOB_LOCATOR and HOMOGENEOUS_BATCH are set to TRUE. java.sql.Types _____ was mixed with java.sql.Types _____.</p>
JZ043	<p>LOB objects are not supported for any of the possible variants of 'ENABLE_BULK_LOAD property', but false. Please consider using alternate setter APIs to insert the data.</p>
JZ044	<p>Server-side locators can not be created within the batch if, SEND_BATCHPARAMS_IMMEDIATE is TRUE. Try using client side LOBs or set SEND_BATCHPARAMS_IMMEDIATE to FALSE.</p>
JZ0BD	<p>Out of range or invalid value used for method parameter.</p>
JZ0BI	<p>setFetchSize: The fetch size should be set with the following restrictions - 0 &lt;= rows &lt;= (maximum number of rows in the ResultSet).</p> <p>Action: Verify that you are calling setFetchSize with the parameter falling within the above range of values.</p>
JZ0BJ	<p>The value set for the IMPLICIT_CURSOR_FETCH_SIZE connection property must be &gt; 0.</p>
JZ0BP	<p>Output parameters are not allowed in Batch Update Statements.</p>
JZ0BR	<p>The cursor is not positioned on a row that supports the _____ method.</p> <p>Action: Do not attempt to call a ResultSet method that is invalid for the current row position.</p>
JZ0BS	<p>Batch Statements not supported.</p> <p>Action: Install or update the SAP jConnect metadata stored procedures on your database with the latest versions.</p>
JZ0BT	<p>The _____ method is not supported for ResultSets of type _____.</p> <p>Action: Do not attempt to call a ResultSet method that is invalid for the type of ResultSet.</p>
JZ0C0	<p>Connection is already closed.</p> <p>Action: Fix the code so that connection object references are nulled when a connection is closed.</p>
JZ0C1	<p>An IOException occurred which closed the connection.</p> <p>Description: The connection cannot be used for any further database activity. If this exception occurs, it can always be found in an exception chain with the JZ006 Exception.</p> <p>Action: Determine the cause of the IOException that disrupted the connection.</p>

SQL state	Message/Description/Action
JZ0CL	You must define the CLASS_LOADER property when using the PRE-LOAD_JARS property.
JZ0D4	<p>Unrecognized protocol in JDBC URL:_____.</p> <p>Description: You specified a connection URL using a protocol other than TDS, which is the only protocol currently supported by SAP jConnect.</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>jdbc::Tds:&lt;host&gt;:&lt;port&gt;</p> <p>If the URL specifies JNDI as a subprotocol, make sure that it starts with:</p> <p>jdbc::jndi:</p>
JZ0D5	<p>Error loading protocol _____.</p> <p>Action: Check the settings for the CLASSPATH system variable.</p>
JZ0D6	Unrecognized version number _____ specified in setVersion. Choose one of the SybDriver.VERSION_* values, and make sure that the version of jConnect that you are using is at or beyond the version you specify.
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>
JZ0EM	<p>End of data.</p> <p>Action: Please report this error to SAP Technical Support.</p>
JZ0F1	<p>SAP ASE 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.</p>
JZ0F2	<p>SAP ASE high-availability failover has occurred. The current transaction is aborted, but the connection is still usable. Retry your transaction.</p> <p>Description: The back-end database server to which you were connected stopped responding, but because you 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>

SQL state	Message/Description/Action
JZ0FP	<p>Incorrect value passed for parameter ____.</p> <p>Description: The value of the parameter specified for the state of the current result set is invalid.</p> <p>Action: Make sure the value is valid: CLOSE_CURRENT_RESULT, KEEP_CURRENT_RESULT, CLOSE_ALL_RESULTS.</p>
JZ0GC	<p>Error casting a ____ as a GSSManager. Please check the value you are setting for the GSSMANAGER_CLASS connection property. The value must be a String that specifies the fully qualified class name of a GSSManager implementation. Or, it must be an Object that extends org.ietf.jgss.GSSManager.</p>
JZ0GK	<p>The ____ array can not be null and has to contain only one key.</p> <p>Description: The autogenerated key column name/index array is either NULL or contains more than one key. Only one key is allowed in the array since it relates to the IDENTITY column.</p> <p>Action: Check the query and correct.</p>
JZ0GN	<p>Error instantiating the class ____ as a GSSManager. The exception was _____. Please check your CLASSPATH and make sure the GSSMANAGER_CLASS property value refers to a fully qualified class name of a GSSManager implementation.</p> <p>Action: Make sure your CLASSPATH environment variable includes any .jar files required by your third-party GSSManager implementation.</p>
JZ0GS	<p>A Generic Security Services API exception occurred. The major error code is _____. The major error message is _____. The minor error code is _____. The minor error message is _____.</p> <p>Action: Examine the major and minor error codes and messages. Check your Kerberos configuration. See <a href="#">Security [page 125]</a>.</p>
JZ0HO	<p>Unable to start thread for event handler; event name = _____.</p> <p>Action: Report this error to SAP Technical Support.</p>
JZ0H1	<p>An event notification was received but the event handler was not found; event name = _____.</p> <p>Action: Report this error to SAP 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>
JZ0I3	<p>Incorrect argument ____ passed to method _____.</p> <p>Action: Verify in the product documentation or in the JDBC API that the correct argument was passed.</p>

SQL state	Message/Description/Action
JZ015	<p>An unrecognized CHARSET property was specified: ____.</p> <p>Action: Enter a valid character-set code for the CHARSET connection property. See <a href="#">jConnect Character-Set Converters [page 47]</a>.</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 SAP jConnect client that supports 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 SAP jConnect version 6.05 or later, and SAP ASE 12.5 or later, you can send your data to the server as <code>unichar</code>/<code>univarchar</code> datatypes.</p>
JZ017	<p>No response from proxy gateway.</p> <p>Description: The connection cannot be established as there is no response from proxy gateway specified with the <code>PROXY</code> connection property.</p> <p>Action: Check the <code>PROXY</code> setting and correct.</p>
JZ018	<p>Proxy gateway connection refused. Gateway response: %ls.</p> <p>Action: Check the proxy gateway settings.</p>
JZ019	<p>This <code>InputStream</code> was closed.</p> <p>Description: You tried to read an <code>InputStream</code> obtained from <code>getAsciiStream</code>, <code>getUnicodeStream</code>, or <code>getBinaryStream</code>, but the <code>InputStream</code> was already closed. The stream may have been closed because you moved to another column or canceled 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 SAP jConnect client that supports all the characters you need to send to the server. You may need to install a different character set on the server, too.</p>
JZ01B	<p>The server's default charset of ____ does not map to an encoding that is available in the client Java environment. Because jConnect will not be able to do client-side conversion, the connection is unusable and is being closed. Try using a later Java version, or try including your Java installation's <code>il8n.jar</code> or <code>charsets.jar</code> file in the classpath.</p>
JZ01R	<p><code>getXXX</code> may not be called on a column after it has been updated in the result set with a <code>java.io.Reader</code>.</p> <p>Action: Remove the <code>getXXX</code> call on the <code>ResultSet</code> column you updated using a reader.</p>

SQL state	Message/Description/Action
JZOIS	<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 these SybResultSet methods: getAsciiStream, getUnicodeStream, getBinaryStream. SAP jConnect does not support this usage.</p> <p>Action: Do not attempt to fetch input streams from columns you are updating.</p>
JZOJO	Offset and/or length values exceed the actual text/image length.
JZOLA	<p>Failed to instantiate Cipher object. Transformation %ls is not implemented by any of the loaded JCE providers.</p> <p>Action: Make sure that the implementation is specified correctly in the JCE_PROVIDER_CLASS connection property of the CLASSPATH.</p>
JZOLC	<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>
JZOMD	<p>ResultSet metadata is not available.</p> <p>Action: Install the metadata stored procedures.</p>
JZONC	<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>
JZONE	<p>Incorrect URL format. URL: _____. Error message: _____.</p> <p>Action: In the URL, make sure that the port number consists only of numeric characters.</p>
JZONF	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.
JZONK	<p>Generated keys are not available because either the Statement.NO_GENERATED_KEYS was used or no keys were automatically generated.</p> <p>Description: The getGeneratedKeys () method cannot return the autogenerated keys because the statement was executed with .NO_GENERATED_KEYS or the statement produced no autogenerated keys.</p> <p>Action: Use getGeneratedKeys () only on statements executed with .RETURN_GENERATED_KEYS, or those that are expected to autogenerate keys.</p>
JZONS	The method ____ is not supported and should not be called.

SQL state	Message/Description/Action
JZOP1	<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 SAP Open Server application, it may indicate an error in the SAP Open Server application that causes the SAP Open Server to send unexpected sequences of results.</p> <p>Action: Use the <code>com..utils.Debug(true, "ALL")</code> debugging tools to try to determine what unexpected result is seen, and to understand its causes.</p>
JZOP4	<p>Protocol error. This message indicates an internal product problem. Report this error to SAP technical support.</p>
JZOP7	<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 re-read 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 be sure that you read columns in ascending column-index order.</p>
JZOP8	<p>The RSMDBA Column Type Name you requested is unknown. This is a SAP internal error; please report it to technical support.</p> <p>Description: SAP jConnect cannot determine the name of a column type in the <code>ResultSetMetaData.getColumnTypeName</code> method.</p> <p>Action: Be sure that your database has the most recent stored procedures for metadata.</p>
JZOP9	<p>A COMPUTE BY query has been detected. That type of result is unsupported and has been canceled.</p> <p>Action: Change your query or stored procedure so it does not use COMPUTE BY.</p>
JZOPA	<p>The query has been canceled and the same response discarded.</p> <p>Action: Check the chain of SQL exceptions and warnings on this and other statements to determine the cause.</p>
JZOPB	<p>The server does not support a requested operation.</p> <p>Description: When SAP jConnect creates a connection with a server, it informs the server of capabilities it wants supported, and the server informs SAP 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 <code>SybConnection.prepareStatement(&lt;sql_stmt&gt;, &lt;dynamic&gt;)</code>, and <code>&lt;dynamic&gt;</code> is set to true, SAP 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
JZOPC	<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 &gt;=6 if you wish to request widetable support.</p> <p>Description: You are trying to execute a statement with a larger number of parameters than the server is configured to handle. The number of parameters that can produce this exception varies, depending on the datatypes of the data you are sending. You never get this exception if you are sending 481 or fewer parameters.</p> <p>Action: Run the query against an SAP Adaptive Server 12.5 or later server. When you connect to the database, set the JCONNECT_VERSION property to "6".</p>
JZOPD	<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 &gt;=6 if you wish to request widetable support.</p> <p>Description: You are trying to execute a dynamic prepared statement with larger number of parameters than the server is configured to handle.</p> <p>Action: Run the query against an SAP Adaptive Server 12.5 or later server. When you connect to the database, set the JCONNECT_VERSION property to "6".</p>
JZOPE	<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 &gt;= 6 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: Run the query against a version 12.5 or later SAP Adaptive Server. When you connect to the database, set the JCONNECT_VERSION property to "6".</p>
JZOPN	<p>Specified port number of ____ was out of range. Port numbers must meet the following conditions: 0&lt;= portNumber &lt;=65535.</p> <p>Action: Check the port number that is specified in the database URL.</p>
JZORO	<p>Result set has already been closed.</p> <p>Action: Fix the code so that ResultSet object references are set to null whenever a result set is closed.</p>

SQL state	Message/Description/Action
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: Verify 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 SAP technical support.</p>
JZOR4	<p>Column does not have a text pointer. It is not a text/image column or the column is NULL.</p> <p>Action: Be sure that you are not trying to update or get a text pointer to a column that does not support text/image data, and verify 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>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>Action: Alter the code so that the application does not attempt to retrieve data from a deleted row.</p>
JZORM	<p><code>refreshRow</code> may not be called after <code>updateRow</code> or <code>deleteRow</code>.</p> <p>Description: After updating a row in the database with <code>SybCursorResult.updateRow</code>, or deleting it with <code>SybCursorResult.deleteRow</code>, you used <code>SybCursorResult.refreshRow</code> 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>
JZOSO	<p>Statement state machine: Statement is BUSY.</p> <p>Description: The only time this error is raised is from the <code>Statement.setCursorname</code> 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 <code>Statement.cancel</code> before setting the cursor name, to make sure that the statement is not busy.</p>
JZOS1	<p>Statement state machine: Trying to FETCH on IDLE statement.</p> <p>Action: Close the statement and open another one.</p>

SQL state	Message/Description/Action
JZ0S2	<p>Statement object has already been closed.</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 <a href="#">Distributed Transaction Management Support [page 113]</a>.</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 <a href="#">Distributed Transaction Management Support [page 113]</a>.</p>
JZ0S8	<p>An escape sequence in a SQL Query was malformed: '_____.</p> <p>Action: Check JDBC documentation for correct syntax.</p>
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: Be sure that each input parameter has a value.</p>
JZ0SB	<p>Parameter index out of range: _____.</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>

SQL state	Message/Description/Action
JZ0SD	<p>No registered parameter found for output parameter.</p> <p>Description: This indicates an application logic error. You attempted to call <code>getXXX</code> or <code>wasNull</code> 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 <code>CallableStatement</code>, that the statement has been executed, and that the output parameters were read.</p>
JZ0SE	<p>Invalid object type specified for <code>setObject()</code>.</p> <p>Description: Illegal type argument passed to <code>PreparedStatement.setObject</code>.</p> <p>Action: Check the JDBC documentation. The argument must be a constant from <code>java.sql.Types</code>.</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 <code>CallableStatement.registerOutParam</code> for more parameters than you declared as <code>OUTPUT</code> parameters in the stored procedure. See <a href="#">RPC Returns Fewer Output Parameters than Registered [page 156]</a>.</p> <p>Action: Check your stored procedures and <code>registerOutParameter</code> calls. Make sure that you have declared all of the appropriate parameters as <code>OUTPUT</code>. Look at the line of code that reads:</p> <pre>create procedure yourproc (@p1 int OUTPUT, ...</pre> <div style="background-color: #f0f0f0; padding: 10px; margin-top: 10px;"> <p><b>Note</b></p> <p>If you receive this error while using SQL Anywhere, upgrade to SQL Anywhere version 5.5.04.</p> </div>
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 <code>oj</code> 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: In addition to following the instructions in the message, also install the latest version of the SAP jConnect metadata.</p>

SQL state	Message/Description/Action
JZOSL	<p>Unsupported SQL type ____.</p> <p>Action: If possible, declare the parameter to be of a type supported by SAP jConnect. Do not use <code>Types.NULL</code> or <code>PreparedStatement.setObject (null)</code>.</p>
JZOSM	<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 <code>JCONNECT_VERSION</code> connection property to a higher value. Or, if possible, try sending your procedure execution command as a language statement.</p>
JZOSN	<p><code>setMaxFieldSize</code>: field size cannot be negative.</p>
JZOSO	<p>Invalid <code>ResultSet</code> concurrency type: ____.</p> <p>Action: Check that your declared concurrency is either <code>ResultSet.CONCUR_READ_ONLY</code> or <code>ResultSet.CONCUR_UPDATABLE</code>.</p>
JZOSP	<p>Invalid <code>ResultSet</code> type: ____.</p> <p>Action: Check that your declared <code>ResultSet</code> type is <code>ResultSet.TYPE_FORWARD_ONLY</code> or <code>ResultSet.TYPE_SCROLL_INSENSITIVE</code>. SAP jConnect does not support the <code>ResultSet.TYPE_SCROLL_SENSITIVE</code> <code>ResultSet</code> type.</p>
JZOSQ	<p>In valid UDT type ____.</p> <p>Description: When calling the <code>DatabaseMetaData.getUDTs</code> method, jConnect throws this exception if the user-defined type is not <code>Types.JAVA_OBJECT</code>, <code>Types.STRUCT</code>, or <code>Types.DISTINCT</code>.</p> <p>Action: Use one of the three UDTs mentioned.</p>
JZOSR	<p><code>setMaxRows</code>: max rows cannot be negative.</p>
JZOSS	<p><code>setQueryTimeout</code>: query timeout cannot be negative.</p>
JZOST	<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 <code>LITERAL_PARAMS</code> connection property is set to false when you execute this query.</p>
JZOSU	<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: SAP ASE and SAP SQL 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>

SQL state	Message/Description/Action
JZ0SV	<p>Combination of setting parameters by name and by index is not allowed in the same CallableStatement.</p> <p>Description: The CallableStatement has parameters specified by name and by index (ordinal position). Mixed use is invalid.</p> <p>Action: Specify parameters by name only or by index (ordinal position) only.</p>
JZ0SW	<p>Invalid ResultSet holdability type: ____.</p> <p>Description: You have specified an invalid value with the <code>setHoldability()</code> method.</p> <p>Action: Use the legal values only: <code>HOLD_CURSORS_OVER_COMMIT</code> or <code>CLOSE_CURSORS_AT_COMMIT</code>.</p>
JZ0T2	<p>Listener thread read error.</p> <p>Action: Check your network communications.</p>
JZ0T3	<p>Read operation timed out.</p> <p>Action: Increase the timeout period by calling <code>Statement.setQueryTimeout</code>.</p>
JZ0T4	<p>Write operation timed out. Timeout in milliseconds: ____.</p> <p>Action: Increase the timeout period by calling <code>Statement.setQueryTimeout</code>.</p>
JZ0T5	<p>Cache used to store responses is full.</p> <p>Action: Use default or larger value for the <code>STREAM_CACHE_SIZE</code> 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 abort, please contact SAP Technical Support.</p>
JZ0T8	<p>Data received for an unknown request. Please report this error to SAP Technical Support.</p>
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>
JZ0TD	<p>Caught ThreadDeath.</p> <p>Description: The calling application thread was killed while SAP jConnect was performing a timed I/O operation.</p> <p>Action: Check the application code to locate the conflict and correct.</p>

SQL state	Message/Description/Action
JZOTE	<p>Attempted conversion between an illegal pair of types. Valid database types are: '_____'.  Description: The database column datatype and the datatype requested in the <code>ResultSet.getXXX</code> call are not implicitly convertible.  Action: Use one of the valid datatypes listed in the error message.</p>
JZOTI	<p><code>jConnect</code> cannot make a meaningful conversion between the database type of _____ and the requested type of _____.  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.  Action: Make sure that the database datatype is implicitly convertible to the object type you want to retrieve.</p>
JZOTO	<p>Read operation timed out.  Description: This exception occurs when there is a socket read timeout.  Action: Increase the timeout period by calling <code>Statement.setQueryTimeout</code>. Also, check the query or stored procedure you are executing to determine why it is taking longer than expected.</p>
JZOTS	<p>Truncation error trying to send _____.  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.  Action: Set the length properly to avoid truncation.</p>
JZOUS	<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>
JZOXC	<p>_____ is an unrecognized transaction coordinator type.  Description: The metadata information indicates that the server supports distributed transactions, but SAP <code>jConnect</code> does not support the protocol being used.  Action: Verify that you have installed the latest metadata scripts. If the error persists, please contact SAP Technical Support.</p>
JZOXS	<p>The server does not support XA-style transactions. Please verify that the transaction feature is enabled and licensed on this server.  Description: The server to which SAP <code>jConnect</code> attempted a connection does not support distributed transactions.  Action: Do not use <code>XADataSource</code> with this server, or upgrade or configure the server for distributed transactions.</p>

SQL state	Message/Description/Action
JZOXU	<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 the user does not have the proper role.</p> <p>Action: Grant the user the role shown in the error message, or have another user with that role conduct the transaction.</p>
JZBK1	<p>SybcBCP class is NOT initialized Please Re-Run MDA sqls to update the MDA stored procedures.</p> <p>Action: Install MDA stored procedures.</p>
JZBK3	<p>Bulk load table does not exists.</p> <p>Action: Correct the table name.</p>
JZBK4	<p>Illegal usage of sql statements mixed with batches in bcp/arrayinsert mode.</p> <p>Description: During a batch operation, you are attempting to execute a nonbatch operation.</p> <p>Action: Wait for the batch operation to complete before attempting a nonbatch operation.</p>
JZBK5	<p>autocommit should be set to true when running bulk load in bcp mode.</p>
JZBK6	<p>Adaptive Server 15.7 or latter and 'allow wide dol rows' DB option must be enabled to insert rows with offsets greater than 8191.</p>
JZBK7	<p>Failed to insert data. Total data size (_____ bytes) exceeds the maximum row size (_____ bytes)allowed for the table _____.</p>
JZBK1	<p>Invalid value set for property ENABLE_BULK_LOAD.</p> <p>Action: Set ENABLE_BULK_LOAD to one of the valid values only – ARRAYINSERT_WITH_MIXED_STATEMENTS, ARRAYINSERT, BCP, or LOG_BCP.</p>
JZNN1	<p>Column does not allow null values.</p> <p>Description: You attempted to set a bit-type column to a NULL value using setNull() in a prepared statement.</p> <p>Action: Examine the query and correct to set a value of 0 or 1 for bit-type columns.</p>
S0022	<p>Invalid column name '_____'. </p> <p>Action: Check the spelling of the column name, and that the named column exists.</p>
ZZ00A	<p>The method _____ has not been completed and should not be called.</p> <p>Action: Check the release bulletin that came with your version of SAP jConnect for further information. You can also check the <a href="#">jConnect Web page</a> to see whether a more recent version of SAP jConnect implements the method. If not, do not use the method.</p>



# 11 Glossary

Glossary of terms used in SAP jConnect for JDBC.

<b>application program interface (API)</b>	a source code based specification intended to be used as an interface by software components to communicate with each other.
<b>SAP Adaptive Server Enterprise (SAP ASE)</b>	manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.
<b>Certicom Security Builder GSE-J</b>	a Java Cryptography Extension (JCE) software cryptographic provider that supports FIPS 140-2 validated cryptographic algorithms.
<b>CyberSafe TrustBroker</b>	a Generic Security Services (GSS) Manager that can be used by jConnect.
<b>Database Server</b>	the back-end system of a database application using client or server architecture.
<b>datatype</b>	a defining attribute that describes the type, values and operations that are legal for a variable.
<b>deadlock</b>	a situation that arises when two users, each having a lock on one piece of data, attempt to acquire a lock on the other's piece of data.
<b>DirectConnect</b>	the ECDA component that provides basic connectivity to non-SAP datasources. In particular, it provides access management, transaction management, and remote systems management through DirectConnect Manager.
<b>distinguished name (DN)</b>	a string that uniquely identifies an entry in the Directory Server. A DN comprises of zero or more relative distinguished name (RDN) components that identify the location of the entry in the directory information tree (DIT). A DN is similar to an analog to an absolute path in a file system in that it specifies both the name and hierarchical location.
<b>GSS library</b>	a library that implements Generic Security Service Application Program Interface (GSS-API).
<b>IETF</b>	Internet Engineering Task Force. The main standards organization for the Internet. A large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet.
<b>SAP jConnect driver</b>	a JDBC driver for SAP servers such as SAP ASE that use Tabular Data Stream (TDS) communication protocol.
<b>Java Cryptography Extension (JCE)</b>	an API that provides a uniform framework for the implementation of security features in Java.

<b>JDBC</b>	Java Database Connectivity. JDBC is a Java API that enables Java programs to execute SQL statements.
<b>JDK</b>	Java Development Kit . An SDK for producing Java programs.
<b>Java Generic Security Services (GSS) Manager</b>	provides a generic interface for authentication and secure messaging.
<b>JNDI</b>	<p>Java Naming and Directory Interface. JNDI enables Java platform-based applications to access multiple naming and directory services.</p> <p>JNDI is an API from Oracle for connecting Java programs to naming and directory services such as DNS, LDAP, and NDS.</p>
<b>Java Runtime Environment (JRE)</b>	part of the Java Development Kit (JDK), a set of programming tools for developing Java applications. May be called Java Runtime.
<b>Java Transaction API (JTA)</b>	an API that allows applications and J2EE servers to access transactions.
<b>Java Transaction Service (JTS)</b>	specifies the implementation of a transaction manager that supports Java Transaction API (JTA) and implements the Java mapping of the Object Management Group Object Transaction Service 1.1 specification at the level below the API.
<b>Java Virtual Machine (JVM)</b>	a virtual machine that provides a platform-independent execution environment that converts Java bytecode into machine language and executes it.
<b>J2EE</b>	Java 2 Platform Enterprise Edition. J2EE is a platform-independent, Java-centric environment for developing, building, and deploying Web-based enterprise applications online.
<b>Kerberos</b>	Kerberos is a secure method for authenticating a request for a service in a computer network.
<b>key distribution center (KDC)</b>	part of a single sign-on (SSO) setup that performs authentication and ticket generation duties.
<b>large object (LOB) datatypes</b>	typically large character objects (text) or binary objects (image).
<b>large object (LOB) locator</b>	contains a logical pointer to LOB data rather than the data itself, reducing the amount of data that passes through the network between Adaptive Server and its clients.
<b>LDAP</b>	Lightweight Directory Access Protocol. LDAP is a software protocol for enabling anyone to locate organizations, individuals, and other resources such as files and devices in a network, whether on the public Internet, or on a corporate intranet.
<b>LDAP Data Interchange Format (LDIF)</b>	a mechanism form representing directory data in text form. The LDIF specification is contained in RFC 2849 and describes a format not only for representing directory data but also a mechanism for making changes to that data.

<b>native-protocol</b>	the native protocol supported by the DBMS to exchange request or response between clients and the server.
<b>net-protocol</b>	a protocol used to exchange request or response between a middle tier gateway that in turn communicates with the database.
<b>object identifier (OID)</b>	an identifier used to name an object. Structurally, an OID consists of a node in a hierarchically-assigned namespace.
<b>primary server</b>	in a high availability (HA) environment, primary server is the server where the client should first attempt to connect.
<b>Replication Server®</b>	maintains replicated data in multiple databases while ensuring the integrity and consistency of the data. It provides clients using databases in the replication system with local data access, thereby reducing load on the network and centralized computer systems.
<b>relative distinguished name (RDN)</b>	a single component within a distinguished name. An RDN comprises of one or more name-value pairs, in which the name and the value are separated by an equal sign (for example, for an RDN of "<uid=ann>", the name is "<uid>" and the value is "<ann>"), and if there are multiple name-value pairs, they should be separated by plus signs (for example, for an RDN of "<cn=Jon Doe+employeeNumber=12345>", the name-value pairs are "<cn=John Doe>" and "<employeeNumber=12345>"). In practice, RDNs containing multiple name-value pairs (called "<multivalued RDNs>") are rare, but they can be useful at times when either there is no unique attribute in the entry or you want to ensure that the entry's DN contains some useful identifying information.
<b>RPC</b>	Remote Procedure Call. A protocol that one program can use to request a service from a program located in another computer in a network without having to understand network details. (A procedure call is also sometimes known as a function call or a subroutine call.) RPC uses the client/server model. The requesting program is a client and the service-providing program is the server. Like a regular or local procedure call, an RPC is a synchronous operation requiring the requesting program to be suspended until the results of the remote procedure are returned. However, the use of lightweight processes or threads that share the same address space allows multiple RPCs to be performed concurrently.
<b>RSA encryption</b>	a highly secure cryptography method.
<b>secondary server</b>	in a high availability (HA) environment, secondary server is the server where client should attempt to connect if connection fails on the primary server.
<b>single-sign-on (SSO)</b>	a session or user authentication process that permits a user to enter one name and password in order to access multiple applications. The process authenticates the user for all the applications they have been given rights to and eliminates further prompts when they switch applications during a particular session.
<b>SAP SQL Anywhere</b>	a fully-featured relational database and data management tool.
<b>SSL</b>	Secure Sockets Layer. SSL is a commonly-used protocol for managing the security of a message transmission on the Internet.



<b>SAP IQ</b>	<p>a high-performance decision-support server designed specifically for data warehousing.</p> <p>SAP IQ is part of the SAP product family that includes SAP SAP ASE and SAP SQL Anywhere. Component Integration Services within SAP IQ provide direct access to relational and nonrelational databases on mainframe, UNIX, or Windows servers.</p>
<b>Tabular Data Stream (TDS)</b>	TDS is an application-level protocol that describes the transmission of data between two computers. TDS defines the types of messages that can be sent, as well as the order in which they are sent. TDS relies on a connection-oriented transport service.
<b>TDS-tunnelling servlet</b>	A servlet that passes through a TDS stream via HTTP or HTTPS packets.
<b>UCS-2</b>	Universal Character Set is an ISO/IEC format for coding character sets. ISO/IEC 10646 was synchronized with Unicode; however, Unicode adds additional constraints, and compliance with 10646 does not guarantee compatibility with Unicode.
<b>UTF-16</b>	Unicode Transformation Format-16 (UTF-16) is a two-byte format in the Unicode coding system.
<b>Wedgetail JCSI</b>	a Generic Security Services (GSS) Manager that can be used by SAP jConnect.

# Important Disclaimers and Legal Information

## Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon  : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

## Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

## Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

## Gender-Related Language

We try not to use gender-specific word forms and formulations. As appropriate for context and readability, SAP may use masculine word forms to refer to all genders.



© 2018 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.